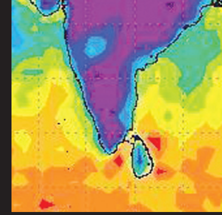

LOS ALAMOS SCIENCE AND TECHNOLOGY MAGAZINE

MARCH 2009

1663

Los Alamos
NATIONAL LABORATORY
EST. 1943

1663

About Our Name: During World War II, all that the outside world knew of Los Alamos and its top-secret laboratory was the mailing address—P. O. Box 1663, Santa Fe, New Mexico. That box number, still part of our address, symbolizes our historic role in the nation's service.

Located on the high mesas of northern New Mexico, Los Alamos National Laboratory was founded in 1943 to build the first atomic bomb. It remains a premier scientific laboratory, dedicated to national security in its broadest sense. The Laboratory is operated by Los Alamos National Security, LLC, for the Department of Energy's National Nuclear Security Administration.

About the Cover: Benno Schoenborn (foreground) and Paul Langan of the Los Alamos Protein Crystallography Station (PCS), where scientists are learning the secrets of how enzymes do their work. The technique they use, neutron protein crystallography, was invented by Schoenborn many years ago, but the PCS has made it practical for scientists worldwide.

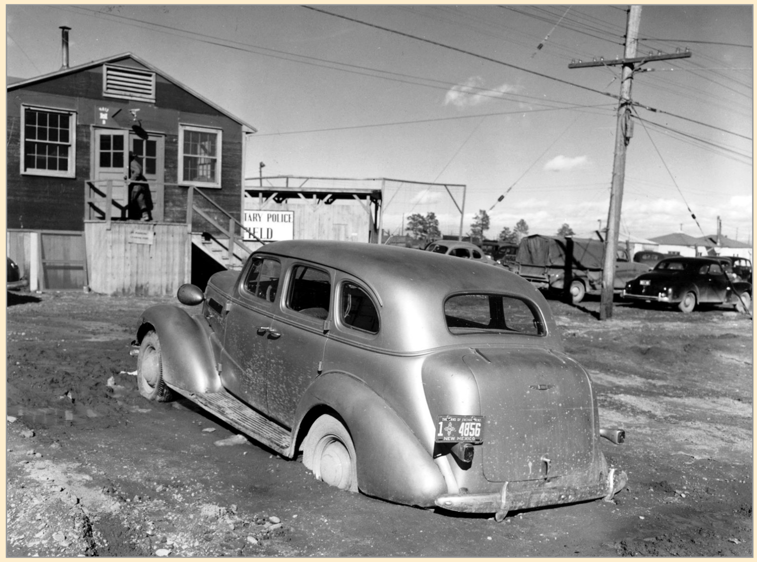

LOS ALAMOS ARCHIVE

During the war years, the absence of pavement often turned Los Alamos roads into muddy quagmires.

From William Priedhorsky

Supporting Science in the 21st Century

America is currently facing energy, security, and environmental challenges that, in their scope and complexity, are perhaps unparalleled in the nation's history. In that context, the national laboratories are being called on to provide the scientific breakthroughs that will be needed to develop long-term solutions.

Scientific breakthroughs, however, seldom arise spontaneously from individual talent, but from a critical mass of talented individuals who are supported by an institution committed to basic research. In addition to facilities and equipment, the institution must provide researchers with financial support to pursue what are often high-risk, big-payoff ideas.

Los Alamos and the other national laboratories not only foster a research environment that is conducive to scientific innovation, but through institutional Laboratory Directed Research and Development (LDRD) programs, also provide critically needed financial support. Congress authorizes the labs to spend at their discretion a small fraction of their budgets (capped at 8 percent for National Nuclear Security Administration labs) in order to build technical capabilities and to explore ways to meet future mission needs.

At Los Alamos, that small financial investment has traditionally yielded an exceptional return. The technical output of LDRD researchers—patent disclosures, peer-reviewed publications, and publications cited by other authors—typically accounts for fully one-quarter of the Laboratory's total. Much of the Laboratory's scientific capabilities, from energy security to large-scale infrastructure modeling, from actinide science to nuclear nonproliferation and detection, can be traced to LDRD investment.

More important, LDRD gives Los Alamos the means to recruit and retain the finest scientific talent. Of the 45 technical staff members hired last year, half were former postdoctoral students supported by LDRD projects. Most of the researchers featured in this issue of 1663 *have received or are currently receiving LDRD funding.*

It is the role of the national laboratories, and especially the national security laboratories, to advance the science that will form the foundation of tomorrow's technology. Through our robust LDRD program, Los Alamos will be able to sustain the scientific workforce required to meet the nation's long-term national security science needs.

Wm. Priedhorsky

Program Director for LDRD

TABLE OF CONTENTS

PetaVision

Simulating the Thoughts behind the Scenes

The Laboratory's Synthetic Visual Cognition Team is using state-of-the-art supercomputers to learn how the brain sees and does other tasks.

The gift of sight is truly amazing. You "instantly" know what everything is in your field of view without seeming to think about it, without asking yourself, "What am I looking at?"

But to give you this knowledge, your brain must quickly make sense of the huge amounts of visual information constantly gathered by your eyes. How does your brain do it?

Scientists don't know exactly.

"Brain research is in a pre-paradigm state," says Garrett Kenyon, a Los Alamos neuroscientist and member of the Laboratory's Synthetic Visual Cognition Team. "We know lots of things about the brain, but we don't really know how it works."

As a result, computer programs designed to emulate the way the brain processes visual information don't begin to approach human levels of performance. For example, an MIT-developed computer-vision program—currently the most-accurate program at identifying objects—misidentifies what it sees 10 percent of the time.

"Imagine that when you crossed the street, 10 percent of the time what you thought was a billboard was actually an oncoming truck," says Luis Bettencourt, leader of the Synthetic Visual Cognition Project. "Clearly this sort of inaccuracy can be lethal in the real world."

So what's missing in the computer programs? What do computers need in order to see as well as people do?

Finding the answer could one day help robots navigate through buildings and cities without running into walls or getting run over and let computers take the wheel of your car in an emergency. It could also allow rapid, automated analysis of the huge volumes of data beamed down each day from reconnaissance satellites or enable computers to identify faces in video

taken at airports—a task at which existing computer methods fail dismally.

Understanding how the brain sees requires a good theory of how the brain works. But neuroscientists disagree about exactly what's needed to formulate such a theory. Research teams all over the world, including the Laboratory's Synthetic Visual Cognition Team, are exploring various possibilities, often aided by advanced supercomputers.

In collaboration with researchers at MIT and elsewhere, the Los Alamos team plans mainly to explore several mechanisms that could improve our understanding of how the brain processes visual information, which should lead to a better understanding of how the brain does all of its tasks. One of the team's major goals is improving the performance of computer-vision software to human levels.

In addition to Laboratory neuroscientists and advanced-computing specialists, the Synthetic Visual Cognition Project features a fairly unique piece of Laboratory hardware—the Roadrunner supercomputer. Roadrunner set the record for supercomputer speed last summer, running software developed by the Synthetic Visual Cognition Team.

Jumping-Off Place

The starting point for several of the team's studies is the MIT computer-vision program. The program implements a model of the primate visual cortex, which is where the brain processes most visual information. The visual cortex is also the best-understood part of the brain. The model is based on experimental studies of how monkeys and humans process visual information.

The team developed a new version of the MIT program to run on "hybrid-architecture" computers such as the Roadrunner supercomputer and a "mini" Roadrunner at the Los Alamos Center for Nonlinear Studies. The new program is called PetANNet, named for the fact that its computer-simulated neurons are connected to compose what's called a neural network, or neural net.

Like the MIT program, PetANNet implements a model of the "what" pathway, the neural pathway that identifies objects in one's field of view. A separate pathway, the "where" pathway, identifies the locations of objects in the visual field. The visual cortex is divided into four major areas from back to front: V1, V2, V4, and IT. The "what" and "where" pathways flow through all four areas, with the "what" pathway on the dorsal side (underside) of the gray matter and the "where" pathway on the ventral (upper) side. In PetANNet, information flows through the "what" pathway almost entirely in the forward direction from V1 to IT, that is, in a "feed-forward" fashion.

Visual information enters the "what" pathway through the lens (cornea) of the eye. The cornea focuses images onto the retina, at the back of the eye, where photoreceptors convert light to the electrical signals the brain's neurons use to communicate with each other. "Roughly speaking," Kenyon says, "your eye has about 500,000 photoreceptors, which is about equal to a half-a-megapixel camera."

The electrical signals from the retina go directly to the back of the brain, to V1, and are then processed through the visual cortex, starting with V1 and ending with IT. The field of view is first characterized in terms of simple visual features present in small square sections of the visual field and then in terms of combinations of simple features that represent more-complex features present in larger sections of the visual field. As the information is processed, individual neurons further up the processing hierarchy recognize features that are more and more complex and present in larger and larger sections of the visual field.

Near the top of the processing hierarchy, in V4, complex features, such as ears and noses, are recognized by individual neurons that view sizable fractions of the visual field—a fact that has been proven through electrophysiology experiments. In IT, individual neurons respond to objects or types of objects that appear anywhere in the visual field, regardless of how they're lit or oriented. "The magic is that an object, say, a face, is identified in IT as belonging to a distinct category regardless of the scene it happens to be part of," says Bettencourt. The V1-to-IT processing hierarchy is illustrated in the figure on the facing page.

Feed-forward processing is thought to determine the minimum time necessary for primates to see and identify objects. Experiments show that when a scene is presented to the visual cortex of a monkey or a human,

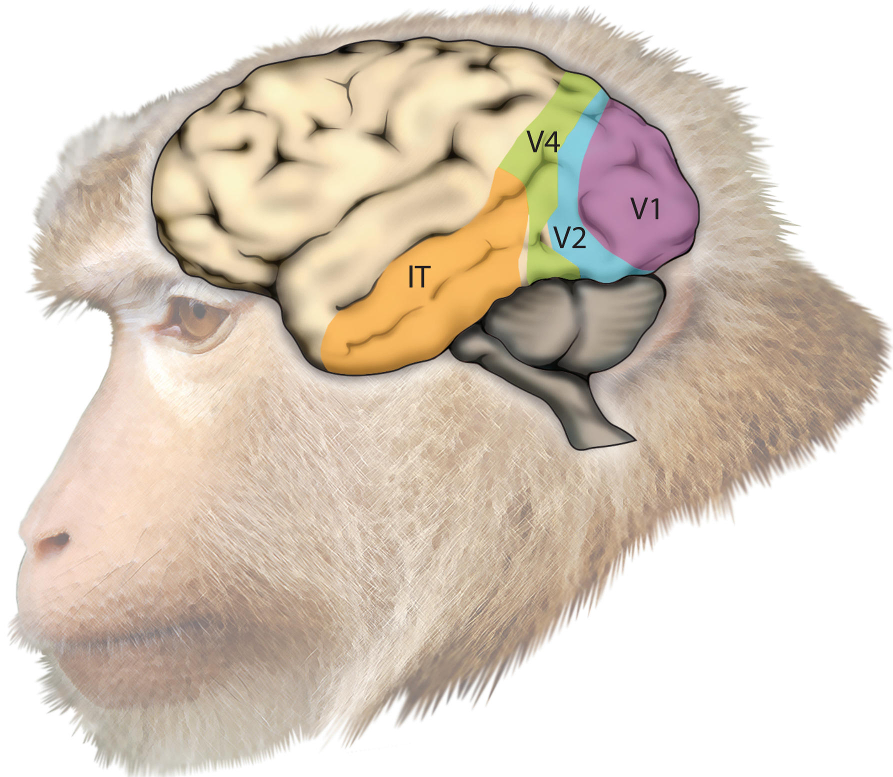

The PetANNet Model

Above: The four major areas in the "what" neural pathway of the primate visual cortex—the pathway that identifies objects in the visual field—are V1, V2, V4, and IT, color coded here to map to the PetANNeT model.

Right: Visual processing begins with V1 dividing the field of view, which in this case contains a dog, into many small square sections, each consisting of about 7 × 7 pixels (photoreceptors). V1 contains "stacks" of neurons, and each stack views one of the small squares. Each neuron in a stack is a "feature detector" that responds to a specific feature in the square viewed by that stack.

The most-prevalent feature detectors in V1 are "edge detectors," each of which is a "simple-cell" neuron sensitive to an edge at a specific angle to the horizontal—exactly horizontal, standing vertical, slanted at 70 degrees, and so on. Other feature detectors are sensitive to, for example, color, spatial frequency, direction of motion, and so on. The visual information next goes to a layer of "complex-cell" neurons, where information is sampled and "pooled." One complex-cell neuron will sample the outputs of several edge detectors sensitive to the same edge angle. By "pooling" information in this way, the complex cell determines if that particular edge angle is present in a larger section of the field of view. Thus begins the process of identifying an object, regardless of where it is in the visual field or how it is oriented or lit.

In V2, a new set of simple-cell neurons monitors the outputs of combinations of the V1 complex-cell neurons over a larger field of view. Each combination represents a new feature that is more complex than those viewed by the stacks in V1 and that is present in a larger swath of the visual field. This process is then repeated in the higher levels of the processing hierarchy, where features are even more complex and appear over even larger swaths.

Finally, in IT, individual neurons are associated with particular objects and categories of objects. Some of these neurons are activated whenever, say, a specific dog appears anywhere in the entire field of view. Others are activated whenever any kind of dog appears.

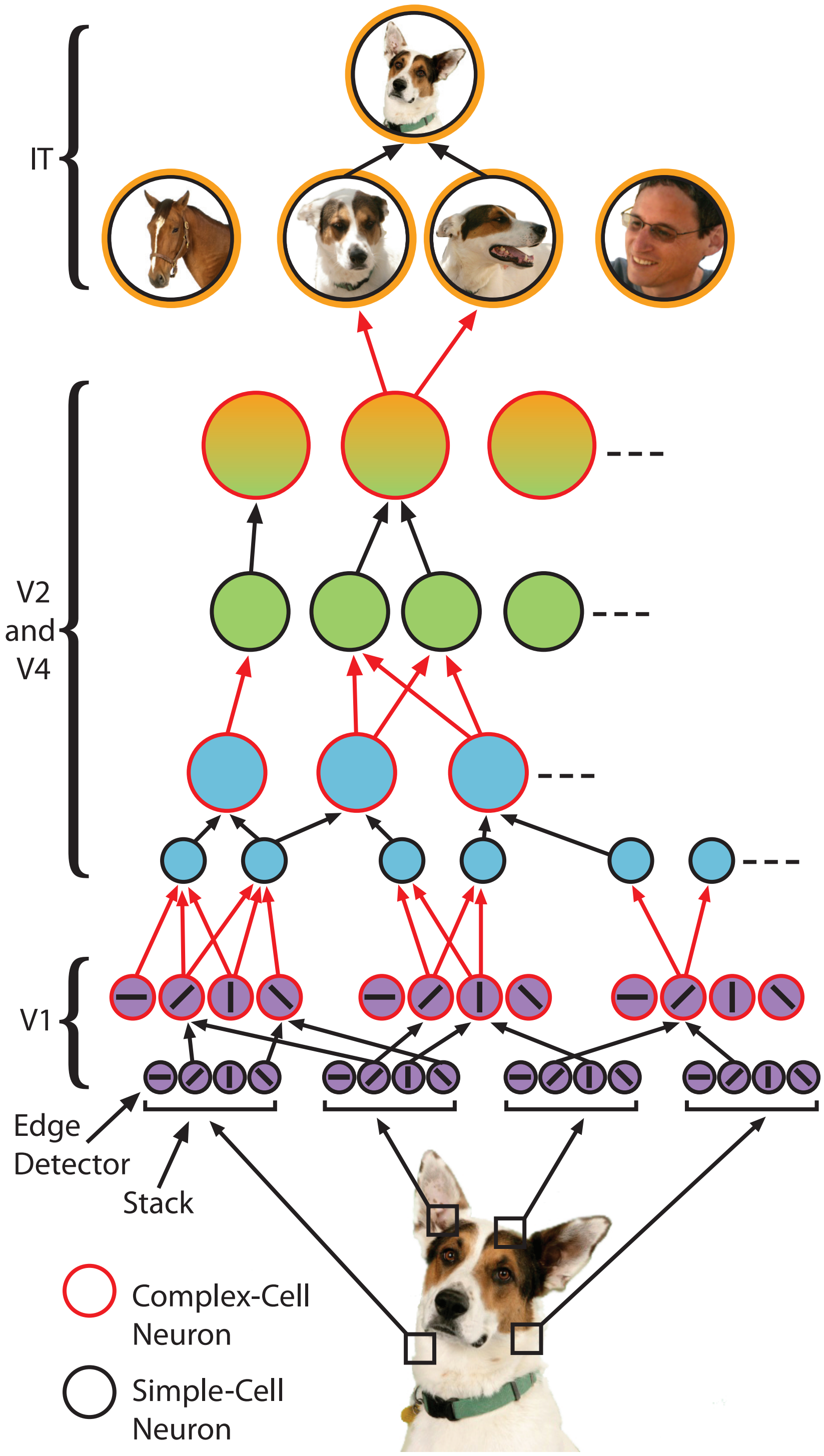

information *initially* flows mainly from the back of the brain to the front—that is, in a feed-forward fashion—rather than laterally or backwards (through "feedback" pathways). The slower processes related to lateral and feedback neural connections kick in *after* the feed-forward processes do, and those connections are not represented in PetANNet. So it's not surprising that when a scene is presented to a human for up to about 50 milliseconds, the human brain identifies objects with about the same accuracy as the program does—70 to 90 percent. But when presented with a scene for longer times, humans become nearly perfect—accurate to at least 99.999 percent. So the question is, how can feed-forward programs be improved?

Grow the Program?

Simply making the program much bigger could help. The feed-forward architecture has roots in the 1950s, when MIT's Marvin Minsky first simulated cortical function by hooking together simulated neurons to form neural nets.

In those days, the limited speed and memory of computers could handle only a small number of neurons and neural connections. Consequently, the neural nets were applied only to very simple problems. The performance of these neural nets was not good, or the problems they solved were trivial. But the proponents of neural nets have claimed ever since that scaling-up the size of the system by adding more neurons to include more feature detectors and more connections would help the simulations learn more about the world and thereby improve their performance to the point that it might eventually rival that of biological cortical material.

"With Roadrunner, we can actually test this hypothesis for the first time," says Bettencourt.

Another member of the Synthetic Visual Cognition Team, Steven Brumby, ran PetANNet on a standard workstation and found that it took about 38 seconds to process a black-and-white image of 320 × 240 pixels. If the model's parameters were scaled up to human values—for example by increasing the number of feature detectors (see illustration at left)—it would take about a day to process a color image of a million pixels, which means that such a simulation could process only 300 or 400 scenes per year! And even if scaling-up significantly improved object-identification accuracy, the software would be much too slow to be useful.

Roadrunner, however, is fast enough to simulate the operation of the entire visual cortex in real time. There are about 10 billion neurons in the human visual cortex, and each neuron is connected to about 10,000 others. Each neuron also fires about 10 times per second, which for a computer means about 10 "floating-point operations" per second (called "flops"). Multiplied together, these numbers give a quadrillion flops per second, or one "petaflop" per second. The speed record the team set with Roadrunner last summer was 1.14 petaflop per second.

So, Roadrunner has what it takes to prove whether scaling-up a feed-forward neural net will improve the software's accuracy to human levels. If scaling-up is the answer, Roadrunner will also be able to identify objects as quickly as humans do.

However, the main research challenge in simulating a system as complex as the visual cortex is teaching the simulation about the visual world. Scaling up means that the representations of the visual world, especially in the upper layers of the visual cortex, can be more numerous and more precise. However, these representations are constructed only when the simulation actually observes the visual world. So, to fully realize the potential for creating more representations that have greater precision, the simulation must also be exposed to the visual world as widely as possible. Thus, the "training set" of visual images used to develop those representations must be as large and diverse as possible.

We also note that humans take months to start seeing well and years to understand what they see. Roadrunner will be able to test new ideas of how the human brain learns about the visual world and how it organizes itself, by making neural connections, to recognize features and to abstract meaning from what it sees.

Plans B

However, team members doubt that scaling-up alone will do the trick. So they are developing other schemes in parallel with the scaling-up approach.

One tack is based on the facts that lateral and feedback neural connections kick in after the feed-forward processes do and that humans identify objects more accurately when scenes are presented to them for at least 50 milliseconds. If the second fact follows from the first, including lateral and feedback connections could improve the model.

In fact, last summer's Roadrunner speed record was set by including lateral connections in a team-developed program called "PetaVision."

PetaVision simulated only area V1—not the entire

COURTESY JUDY HEDDING, ABOUT.COM:PHOENIX

"what" pathway. Nor did PetaVision include feedback connections. So, PetaVision couldn't test whether lateral and feedback connections can together or separately improve the performance of the entire visual-cortex model. However, PetaVision did show that lateral connections can be important in processing visual information, and the code paved the way for testing the effects of lateral connections in models that include all four major areas of the visual cortex. It also tested some other promising approaches.

PetaVision's neurons were edge detectors whose lateral connections to other neurons were "weighted" to detect smooth curves in the simulated visual field; the weighting was derived from the results of experiments. First, as in PetANNet, each small square of pixels in the visual field was analyzed by a stack of feature detectors—in this case, only edge detectors (see illustration on previous page). However, unlike in the feed-forward model, a PetaVision neuron that detected its targeted edge orientation sent lateral signals to other neurons. If a nearby neuron detected an edge that made a small angle with the edge detected by the first neuron, the weighting caused both neurons to send more signals to each other, generating a local feeding frenzy of neural activity. In this way, the neurons corresponding to segments of a smooth curve became highly active, while neurons corresponding to squares that were blank or contained edges with comparatively large angles were suppressed and became listless.

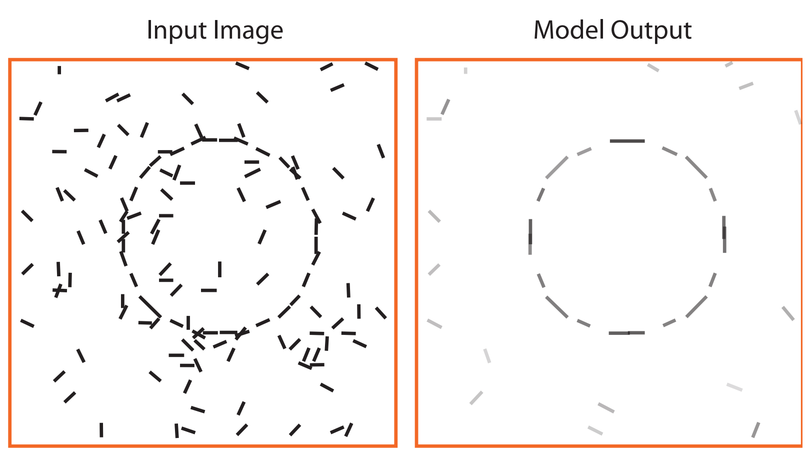

The "Input Image" at right was input to the PetaVision program, which was run on Roadrunner last summer. Using only the simple-cell and complex-cell neurons described on pp. 4 and 5, with lateral connections between the simple-cell neurons that were weighted to emphasize smooth curves, PetaVision found the circle embedded in the input image, as shown in the "Model Output." Such "segmentation" tasks are key to determining the positions of objects in the visual field.

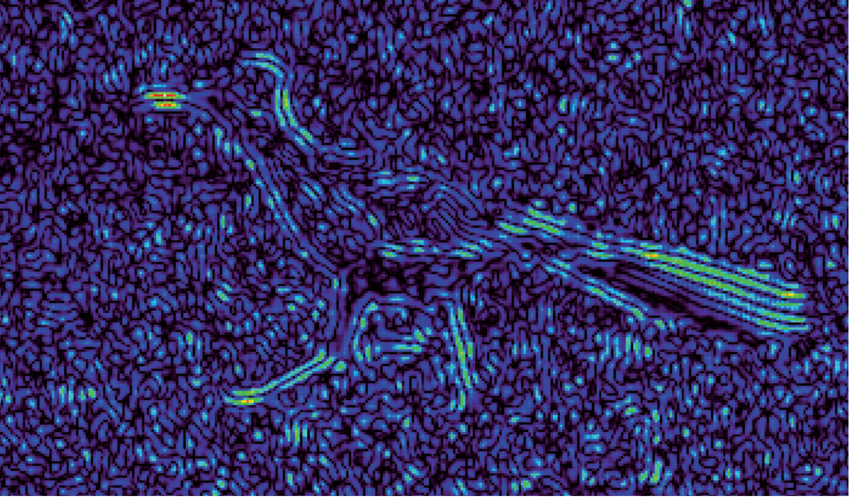

Left: A photo of a roadrunner, the New Mexico state bird, was used as input to the PetANNet program, running on Los Alamos' Roadrunner supercomputer. Right: An interpretation of what PetANNet's complex-cell neurons "saw" in response to the input photo.

This weighting of lateral connections allowed PetaVision to do something PetANNet could never do in its present state: find the border of a circle (see figure on the bottom of the facing page). This may not seem particularly earthshattering, but finding the borders of an object in one's visual field—which is called "segmentation"—is essential to identifying the object's location. Simulations of this sort could pave the way for exploring the more poorly understood "where" pathway of the visual cortex.

Bettencourt also points out that PetaVision found the circle using a *prescribed* weighting. However, the team plans to modify the software so the neurons learn how to find smooth contours *on their own*, just as we do.

Finally, in contrast to the very-simple neuron model in PetANNet, the electrical signals sent between PetaVision's neurons—and the neurons' responses to those signals—were modeled in biological detail.

Biological neurons talk to each other by sending out impulses, "spikes," of voltage. Each spike lasts about a millisecond. PetaVision modeled each spike's amplitude and duration, along with the spike's precise placement in time.

Precise spike timing is known to be used by the cortical tissue that processes auditory information, for example, in bats, "who are geniuses of sound," Bettencourt says. Some of the neurons in the auditory cortical tissue of bats locate sounds by measuring the difference between the placements in time of two spikes to a precision as small as 10 percent of a spike's duration. (The distance from a source of sound is usually slightly different for each ear, so the associated neural signals are slightly displaced from each other in time.) Moreover, Kenyon has studied spike timing in cat retinas, where its importance has also been shown. PetaVision's accurate spike-timing model could help the Synthetic Visual Cognition Team see if spike timing could be important in other cortical activities as well.

Playing Off Each Other

Both Roadrunner and the brain quickly and efficiently process huge amounts of information. There are striking similarities—and differences—in how they do so.

For example, each of Roadrunner's microprocessors performs about one billion operations per second, whereas a neuron performs about a thousand operations per second. However, Roadrunner—even though it is a "green" supercomputer—consumes about 2.341 megawatts of power, enough to run two thousand homes. (The imposing stack of Roadrunner's giant cooling towers, which dissipate the huge amounts of heat generated by the supercomputer's thousands of superfast chips, is a distinctive feature near the Los Alamos building that Roadrunner calls home.) However, because the neurons in the brain operate much more slowly than do a supercomputer's microprocessors and because the brain is far more parallel than a supercomputer is, the brain uses only 20 to 30 watts!

As research programs such as the Synthetic Visual Cognition Project help us learn how cortical circuits work, we may one day be able to build hardware that can do what the brain does with much less power than existing supercomputers need—or that can operate much faster than existing brains do! Meanwhile, PetANNet and PetaVision are proof that the computational limitations of supercomputers are no longer major obstacles to studying the brain as an integrated system.❖

—Brian Fishbine

Garrett Kenyon (left) and Luis Bettencourt (right) lead the two major projects at Los Alamos that deal with simulating brain function. Kenyon heads a project funded by the National Science Foundation to develop a "neural workbench" that can be used to simulate various kinds of cortical tissue. Bettencourt heads the Synthetic Visual Cognition Team, which is funded by the Laboratory Directed Research and Development Office and which focuses on simulating the cortical tissue that processes visual information.

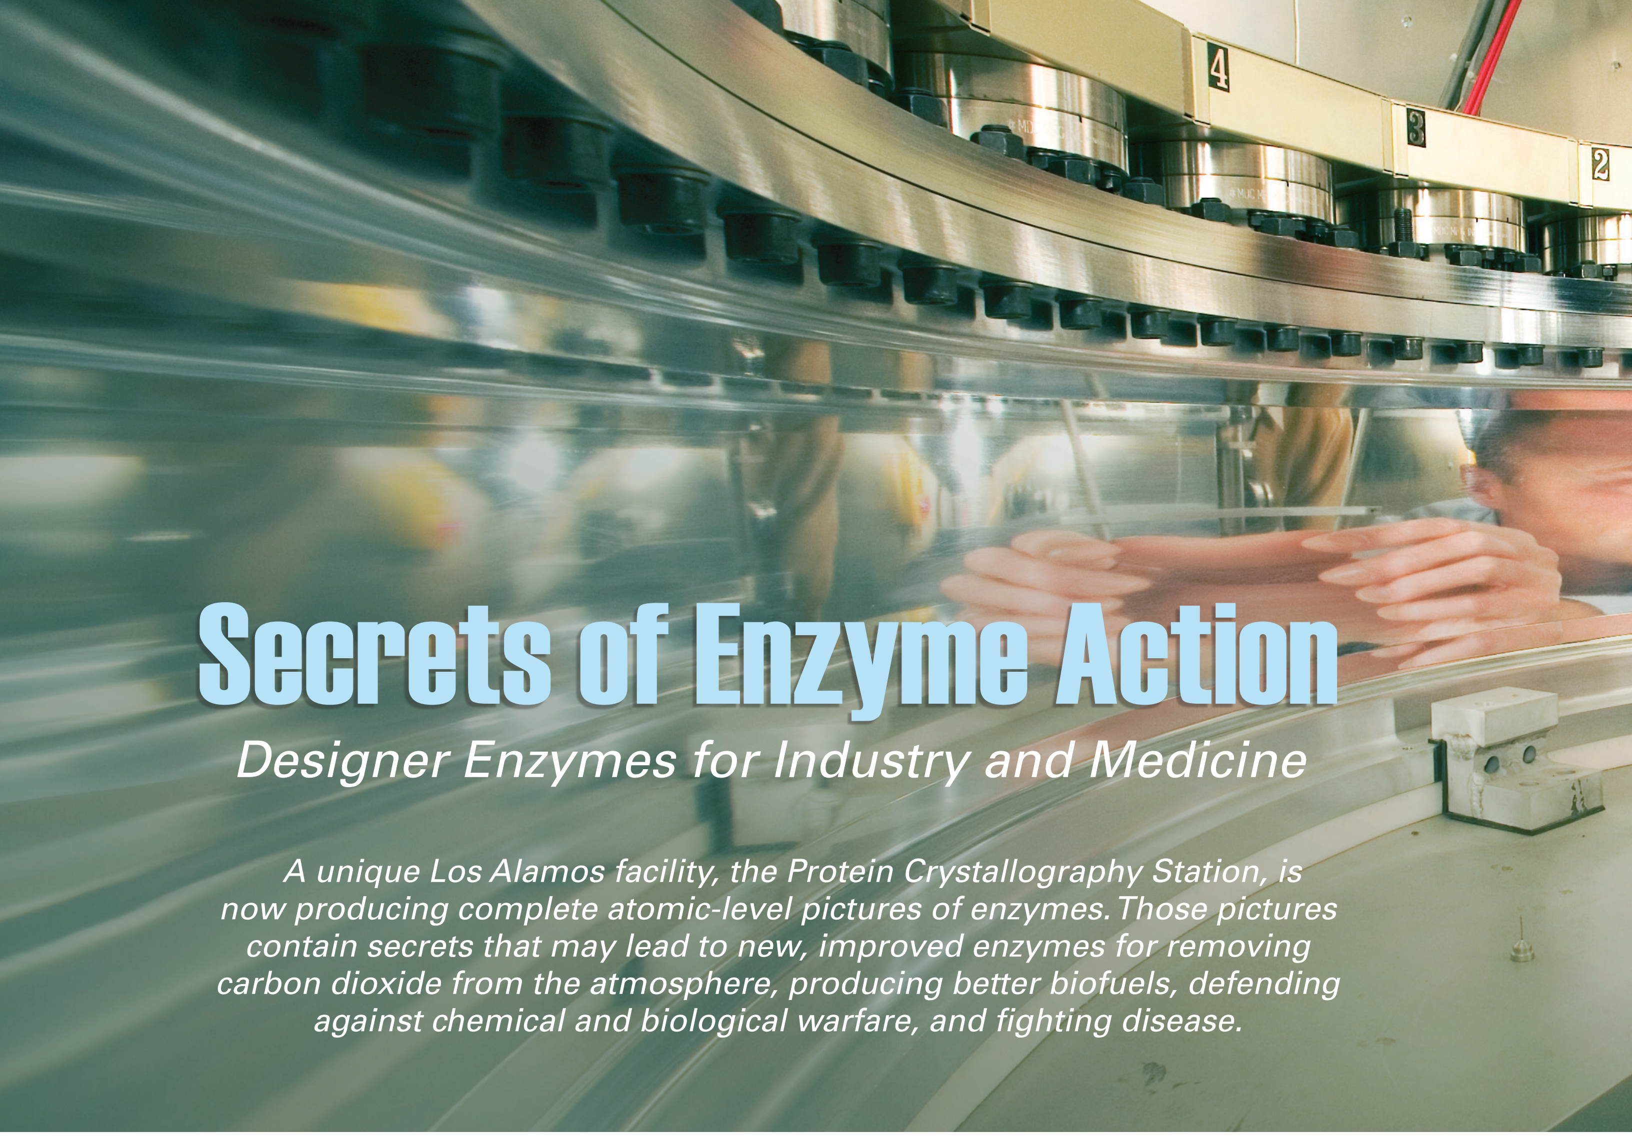

Secrets of Enzyme Action

Designer Enzymes for Industry and Medicine

A unique Los Alamos facility, the Protein Crystallography Station, is now producing complete atomic-level pictures of enzymes. Those pictures contain secrets that may lead to new, improved enzymes for removing carbon dioxide from the atmosphere, producing better biofuels, defending against chemical and biological warfare, and fighting disease.

Andrey Kovalevsky, postdoctoral fellow at the Los Alamos Protein Crystallography Station (PCS), stares intently at a computer screen filled with dark dots on a white background. It looks like the old photographic plates of stars in the night sky. But Kovalevsky is not looking for astronomical objects. He's trying to infer the positions of atoms in a crystallized enzyme. He hopes to uncover the complete structure and operation of xylose isomerase, an important enzyme that catalyzes the conversion of one type of sugar into another.

Enzymes are large protein molecules made by living organisms. They are the most-powerful catalysts known on Earth, speeding up biochemical reactions that, in their absence, would never occur. How is that possible? At the PCS, Kovalevsky and other young scientists hope to learn the enzymes' secret ways and then re-engineer them to do new tasks or perform old ones more efficiently.

Until the PCS opened 5 years ago at the Los Alamos Neutron Science Center (LANSCE), most scientists used only x-ray crystallography to learn the atomic-level structure of proteins. Researchers would shine an x-ray beam on a crystallized protein and record the diffraction "peaks"—intense spots created when x-rays diffract (scatter coherently) from the crystal's orderly array of atoms. From the pattern and intensity of those peaks, scientists could deduce the three-dimensional arrangement of atoms making up the protein's structure.

An astonishing 60,000 proteins have been analyzed this way, but for enzymes, the resulting structural models have one serious limitation—the hydrogen atoms are missing. X-ray diffraction peaks from hydrogen are usually too faint to see.

Diffraction peaks are created when x-rays scatter from the electron cloud surrounding an atomic nucleus. Hydrogen has only one electron in its cloud, so it scatters x-rays very weakly. With 50 percent of the atoms in an enzyme being hydrogen, x-ray crystallography can provide only an incomplete picture

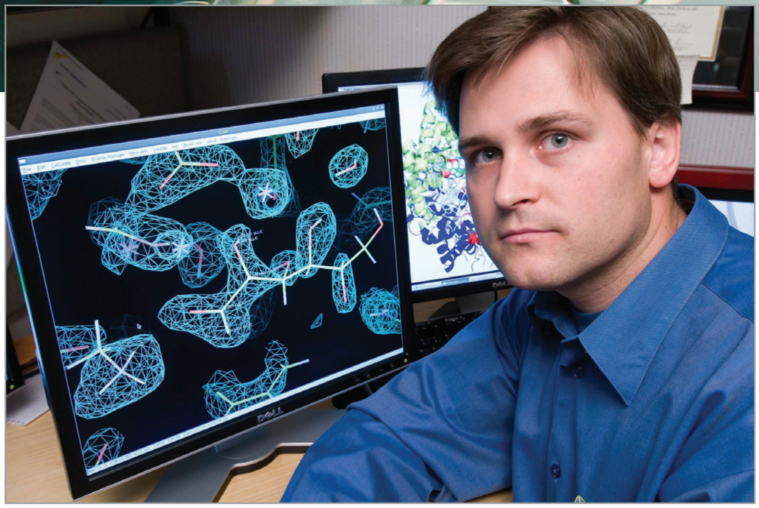

of an enzyme's structure.

Marc-Michael Blum, a German researcher using the PCS, explains, "Some enzymes are like Swiss army knives. Their surfaces look almost alike, but they contain very different tools. Only by knowing the atomic arrangement of those tools, including the positions of the hydrogen atoms, can you figure out how an enzyme really works."

The positions of the hydrogen atoms are especially important because they invariably get shuffled about during enzyme-catalyzed reactions. The hydrogen positions give critical cues about how the reaction moves forward.

"Researchers now have a way around the hydrogen problem," says Paul Langan, team leader of the PCS. "Here at the PCS, they can use neutrons, not x-rays, to do the crystallography and within a few weeks have enough data to locate the hydrogen positions relative to the other atoms in an enzyme." (See "Birth of a New Technique.")

Unlike x-rays, neutrons scatter as strongly from hydrogen as from other elements because they scatter from the atomic nucleus rather than the electron cloud. But the scattering strength depends on the nuclear composition. By replacing a crystal's hydrogen with the isotope deuterium, researchers can add a distinct

Above: Paul Langan (foreground) and Benno Schoenborn aligning the beamstop that protects the gleaming neutron detector of LANSCE's Protein Crystallography Station. The detector, developed by Brookhaven National Laboratory, is a major advance and makes the PCS a practical tool. Inset: Andrey Kovalevsky working on an atomic-level model of the enzyme xylose isomerase. PHOTO BY DIXON WOLF

DFPase, an enzyme that destroys nerve agents, was discovered accidentally in squid during early experiments on nerve conduction.

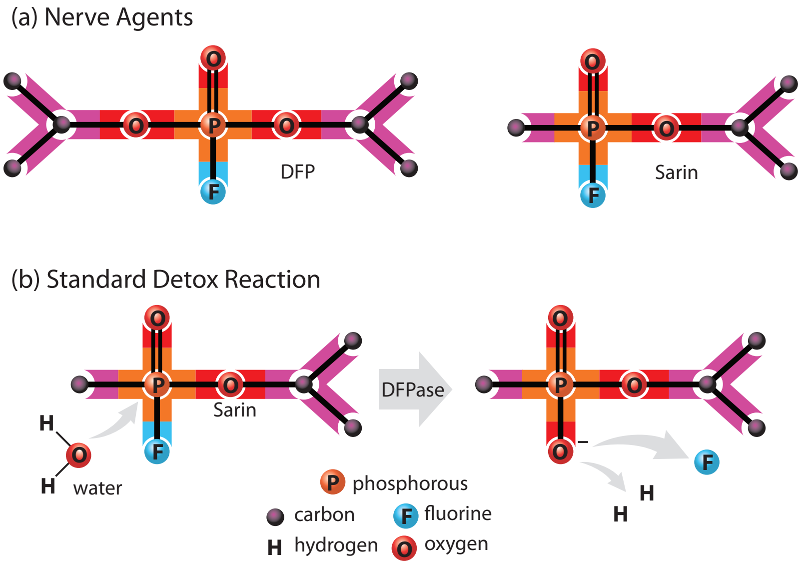

signal component that helps them locate the positions of the replaced atoms. Conveniently, Mary Jo Waltman, a crackerjack technologist, is available to help PCS users grow deuterated crystals in special deuteration laboratories.

Kovalevsky adds, “If we can crystallize the enzyme at various catalytic stages, we can visualize, step by step, how it changes shape and shuffles hydrogens to different locations.”

Seeing this level of detail has reveled a big surprise: some enzymes take a much more active role in the reaction chemistry than commonly thought.

That conclusion is based on x-ray structures, the results of biochemical experiments in solution, theoretical calculations, and more. “The neutron work merely ‘dots the i’s,’” comments Benno Schoenborn, inventor of neutron protein crystallography. “But that final level of information keeps the theorists honest.”

It can also provide a firm basis for engineering an improved version of an enzyme, as a group of German researchers is showing.

A New Defense against Nerve Agents

For many years now, a German team has been studying an obscure enzyme found in squid. Why? Because the enzyme could serve as a defense against nerve agents. Called DFPase because it catalyzes the destruction of the nerve agent DFP (di-isopropyl fluorophosphate), this enzyme is effective against a number of phosphorus-based nerve agents, including Sarin, the substance that killed 12 and sickened thousands in the Tokyo subway attacks of 1995.

Two members of the team, Blum, who has been working with the medical branch of the German army, and Julian Chen, a California transplant and an assistant professor of biophysical chemistry at Goethe University in Frankfurt, recently came to the PCS to clinch their findings on the inner workings of DFPase. The new findings, published early this year in the *Proceedings of the National Academy of Sciences* (PNAS), not only have overthrown past ideas about how the enzyme works, but also have led to an engineered version of the enzyme, speeding its activity rate to the point that it can safely decontaminate exposed skin and sensitive surfaces on optics and electronics in less than 10 minutes. The engineered enzyme also works on another nerve agent, VX, which was immune to the action of DFPase.

All these nerve agents act by quickly blocking the nervous system’s “off switch” for muscles and glands, leaving them stuck in a semi-active state. As a result, they may tire, and in the case of large doses, the body can lose the ability to sustain breathing.

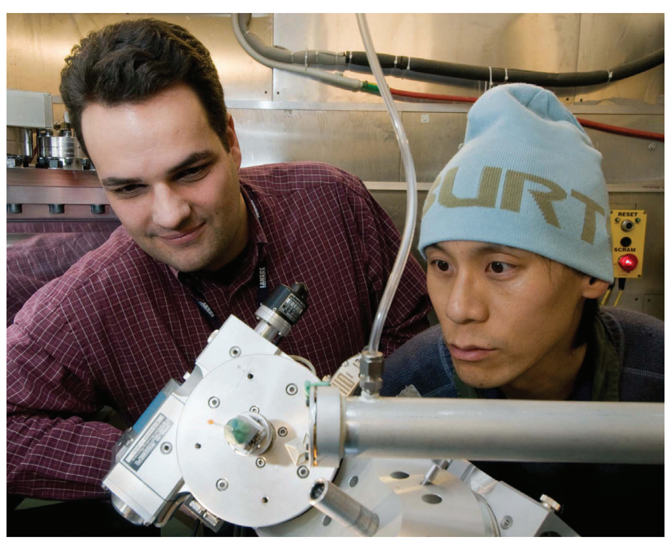

Marc-Michael Blum (left) and Julian Chen examining the PCS’s sample holder. PHOTO BY DIXON WOLF

Birth of a New Technique

Amidst all the excitement and bustle at the Los Alamos Protein Crystallography Station (PCS) is an intense, urbane gentleman, Benno Schoenborn, giving very-specific advice on how to get the most out of the instrument he and his colleagues have designed. Schoenborn first dreamt of making neutron crystallography a tool for biology more than 40 years ago.

Back then, neutrons were available only at nuclear reactors, and safe reactor operations put a firm lid on neutron intensity (the number of neutrons per second, per square centimeter, coming out of the reactor). Although Schoenborn used reactor neutrons to prove that neutron crystallography would work on proteins, he saw it as a labor of love, taking many, many months to collect enough data for a single protein structure.

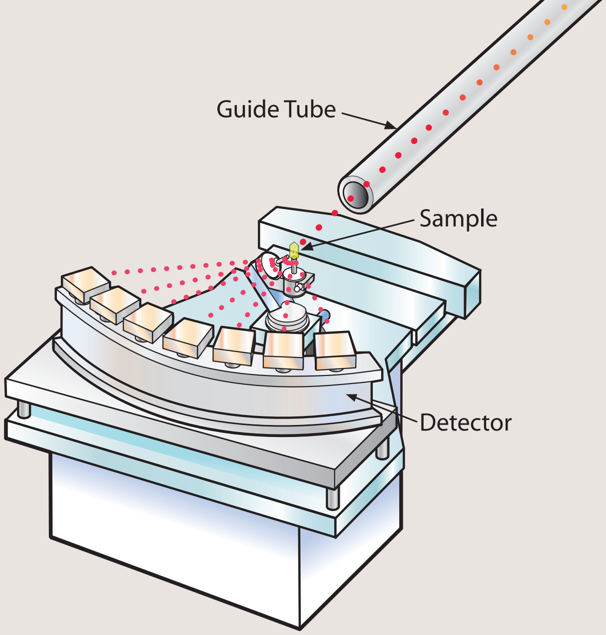

A neutron pulse striking and diffracting from a sample at the PCS.

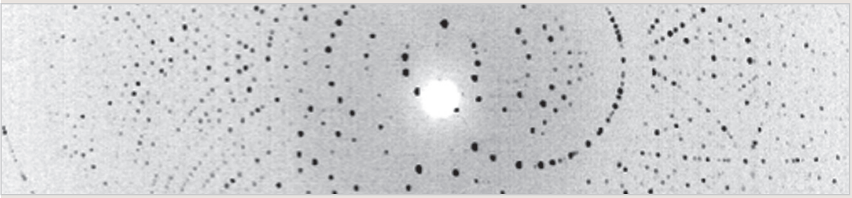

A neutron diffraction pattern.

The intensity lid came off with the advent of the spallation neutron source at the Los Alamos Neutron Science Center (LANSCE). A high-intensity pulse of neutrons is produced each time a proton pulse from the LANSCE accelerator hits a tungsten target. Some of those neutrons are directed down an evacuated tube to the crystal mounted in front of the PCS neutron detector.

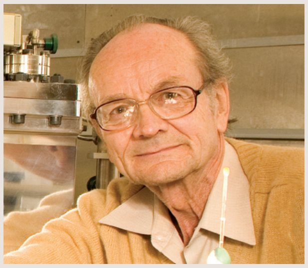

Benno Schoenborn

Like runners coming out of the gate, the neutrons stretch out according to speed along an evacuated guide tube, the fastest arriving at the crystal sample first and the slowest 10 milliseconds (1/1000th of a second) later. Each neutron's arrival is like a time stamp that announces its speed, or energy. Neutrons of different energies are diffracted by different sets of planes in the crystal. The PCS records the diffraction patterns created at 100 different energies (times) within the duration of each 10-millisecond pulse.

Schoenborn, with colleagues Eric Pitcher, Phil Ferguson, and later Paul Langan, designed every component of the PCS to maximize the neutron intensity and reduce background scattering so scientists could collect excellent data in 15 to 30 days of beam time using much-smaller crystals. The big payoff is that the larger proteins, for which large crystals are nearly impossible to grow, have become eligible for neutron crystallography.

The original belief was that DFPase did its defensive work by catalyzing a reaction between water and the nerve agent. As shown in the diagram on the facing page, the standard picture was that oxygen from the water attacked by binding to the phosphorus in the nerve agent. That attack caused the fluorine that is bound to the phosphorus to be jettisoned, thereby detoxifying the nerve agent. (Without the fluorine, the agent no longer binds strongly to the critical off switch.) This detox reaction was presumably catalyzed in the area of the enzyme in which chemical reactions take place, DFPase's active site.

Before testing that theory, DFPase research was focused on just isolating the enzyme, that is, getting enough of it to work on. "At one point my Ph.D. advisor's group came back from the fish market with a couple hundred squid," says Blum. "We needed that many just to get micrograms of the enzyme." Fortunately, that amount was enough to determine the unique sequence of 314 amino acids that form the precious enzyme.

With the sequence in hand, the researchers could exploit gene technology to manufacture the enzyme in much larger quantities and pursue a slew of different

Diagram, facing page: (a) DFP, Sarin, and other phosphorus-based nerve agents typically have a fluorine atom (blue) bound to a phosphorus atom. (b) It was thought that DFPase catalyzed a reaction in which the fluorine was jettisoned and an oxygen from a water molecule swooped in to replace it. The research reported here shows otherwise.

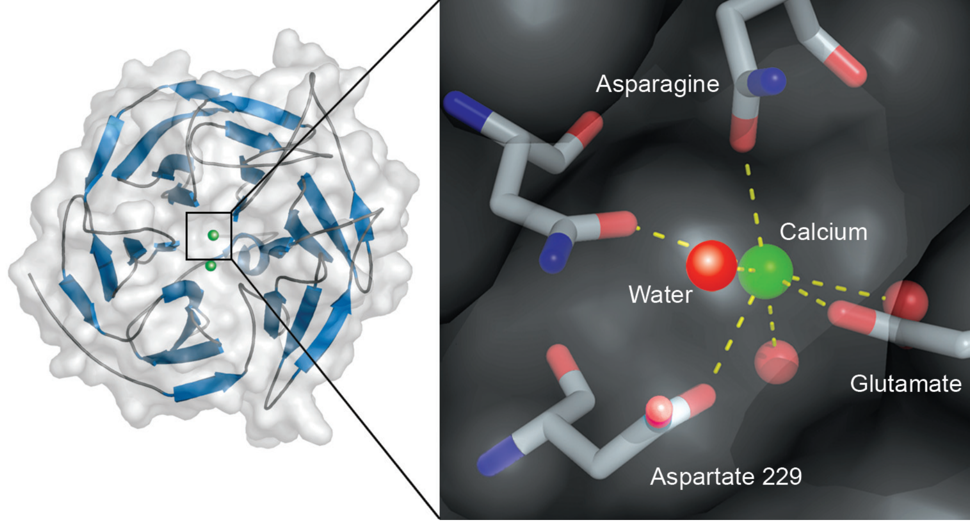

Top view of the enzyme DFPase (white) shows six sections (the groupings of blue ribbons), arranged in a circle around a central tunnel containing two calcium ions (green). The central calcium holds the enzyme together; the upper calcium is in the enzyme's active site. Blowup: In the x-ray structure of DFPase's active site, four amino acids and three oxygen atoms (red) —presumably water molecules—surround and are bound to the calcium ion. Without this calcium ion, the enzyme is not active against nerve agents.

techniques to probe its mode of operation.

X-ray crystallography revealed the structure of the enzyme (minus the hydrogen atoms, of course). The active site, where the enzyme presumably holds the nerve agent and the water in close proximity, turned out to be an indentation at the top where four of the enzyme's amino acids and three isolated oxygen atoms (red) were all loosely bound to a central calcium ion (see blowup in the figure above). The oxygen atoms were presumably water molecules (with their hydrogen atoms invisible) that had been frozen into position during the crystallization of the enzyme in the presence of water.

Blum says, "It was not clear which amino acid or water molecule would act to attack a nerve agent's phosphorus atom."

To uncover the active player, the team co-crystallized the enzyme with a nerve agent surrogate and again used x-rays to determine the combined structure.

To everyone's surprise, the nerve agent surrogate pushed the topmost water molecule out of the active site rather than interacting with it. Moreover, the surrogate had bound to the calcium in an unexpected and very exacting configuration. It looked as if one of the amino acids in the active site, rather than water, could easily provide the oxygen atom for detoxifying the nerve agent.

Blum and Chen began to suspect that, contrary to previous assumptions, a catalyzed reaction with water is not what destroys a real nerve agent. They proposed, instead, that a negatively charged oxygen from the enzyme's amino acid aspartate 229, is what binds to the nerve agent, causing the ejection of a bound fluorine atom (see figure at right). A water molecule later replaces the enzyme's oxygen.

To verify this radical proposal, the researchers did a series of tracer experiments in which the enzyme and the nerve agent DFP were surrounded by water in which oxygen-16, the common oxygen isotope, had been replaced by the heavier oxygen-18. By showing that the oxygen-18 atoms in the water showed up in the reaction *not initially* but only *after* a given DFPase molecule had acted multiple times, Blum and Chen were able to demonstrate that DFPase was indeed the source of oxygen that detoxified the nerve agent.

But the critics remained unconvinced, claiming that an invisible hydrogen ion was present on the aspartate 229's oxygen, neutralizing its negative charge and thereby preventing it from making a successful attack. These critics also claimed that the necessary oxygen came instead from a hydroxide ion bound to the calcium in the enzyme's active site.

To find out who was right, Blum and Chen abandoned x-ray crystallography and mounted a PCS neutron crystallography experiment to pinpoint the hydrogen atoms' positions in the DFPase structure.

Chen came to Los Alamos from Frankfurt to grow

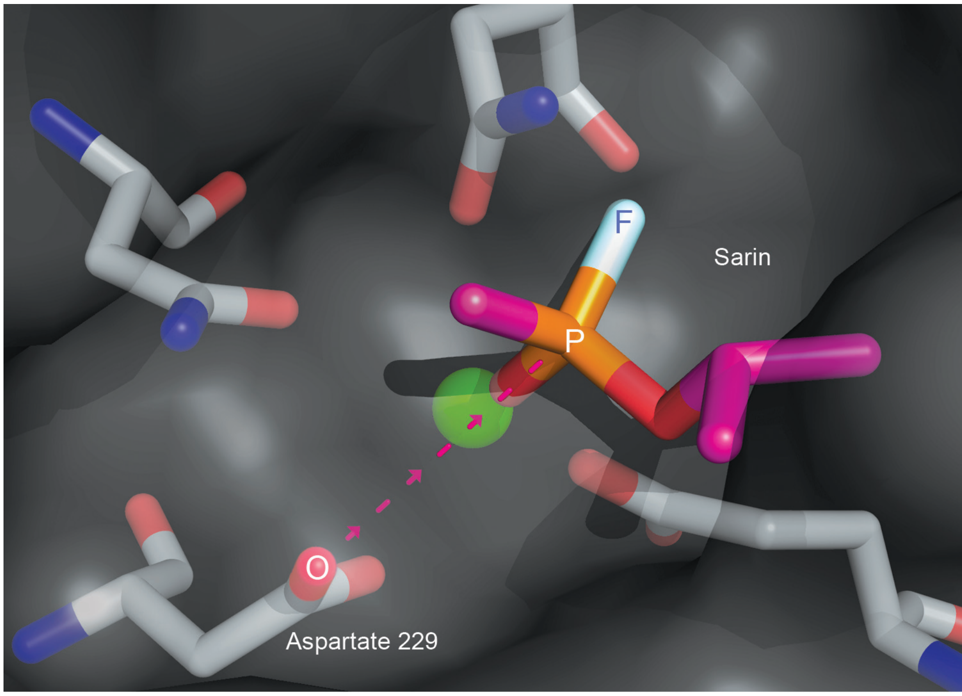

Blum and Chen's newly discovered detox mechanism. An oxygen from aspartate 229 lines up with the fluorine-phosphorus chemical bond on the nerve agent Sarin. It then attaches to the phosphorus and jettisons the fluorine (not shown), destroying the nerve agent's toxicity.

the needed crystals of DFPase and some weeks later returned for the delicate task of mounting a large, soft 2.4-mm-long crystal in a capillary tube. Enzyme crystals are typically soft and delicate because they are about 50 percent water.

Blum and Chen used this crystal to record 37 neutron diffraction images during one month. They then took advantage of new computer programs developed by Langan, Marat Mustyakimov (staff member at Los Alamos), and colleagues from Lawrence Berkeley Laboratory to compare x-ray diffraction data with the new neutron diffraction data and determine the most-likely structure for DFPase crystallized in the presence of water.

The neutron work unequivocally showed that the topmost oxygen atom in DFPase's active site had two hydrogen ions bound to it. It was indeed water, not a hydroxide ion. Also, there was no hydrogen ion on aspartate 229's active oxygen; it was negatively charged. All objections to the new theory about the detox mechanism had been knocked down.

"Our practical goal was to design better versions of DFPase," says Blum. "With proof for the new detox mechanism in hand, we redesigned the shape of the enzyme's active site so that the most-toxic versions of each nerve agent would naturally bind in the best orientation for an aspartate 229 attack. Our best new version of DFPase worked 2 to 10 times faster on all the known nerve agents than the original enzyme did, and it will likely become an important defense for first responders in the case of a nerve agent attack."

In the Wings

The DFPase success story may soon be repeated with other enzymes being studied at the PCS.

Kovalevsky is focused on improving the effectiveness of xylose isomerase (XI) because it's used in the food industry to convert glucose from starch into the much-sweeter fructose. XI is also used in the biofuel industry to convert xylose, a sugar derived from woody plants, into xylulose, a sugar that microorganisms more readily convert into ethanol.

XI has been studied for 30 years, but it took the PCS to locate the hydrogens. Kovalevsky has found them not only in the native enzyme but also in the enzyme crystallized with an intermediate form of the sugar. An unexpected amino acid, far from the purported active site, had removed a hydrogen from the sugar, not donated one, as expected for this type of reaction. The reaction seems to be proceeding in new way, one first intuited from x-ray and solution experiments. Kovalevsky needs neutron diffraction results from two more stages of the reaction to pin down the exact mechanism.

Zoe Fisher, with her baby Owen, asks Marat Mustyakimov to admire the model of carbonic anhydrase, her favorite enzyme.

Also under study at the PCS is a type of carbonic anhydrase (CA) found in humans—HCAII—the fastest acting of all the enzymes that catalyze the conversion of carbon dioxide to bicarbonate. HCAII is being extensively developed for potential use in the sequestration of carbon—capturing and chemically changing carbon dioxide for underground burial. The rate of catalysis is limited by how fast the enzyme moves hydrogen ions away from its active site. Using recent PCS neutron data, Los Alamos postdoctoral fellow Zoe Fisher and researchers from the University of Florida hope to map out this transfer pathway, information that could lead to improving this enzyme's performance.

Impact on the Biological Community

At this point the PCS is a unique instrument with many more subscribers than can be handled during the year.

Schoenborn believes this is just the beginning for neutron crystallography. "We've only scratched the surface. The information we gather will increase in value as scientists, especially the drug design people, need more details about enzyme action."

Langan agrees. "The greater the demand for new, improved enzymes to address new challenges in renewable energy, the environment, chemical and biological threat reduction, and therapeutics, the more scientists will turn for answers to neutron crystallography and the greater the need will be for new instruments like the PCS." ❖

—Necia Grant Cooper

The Cloudy Science of Aerosols

Seeking the Unknown Quantity in Climate-Change Predictions

Models of climate change will fail to provide accurate predictions unless they account for the impact of aerosol particles and the clouds that form around them. Los Alamos researchers have taken to the skies to quantify the effects of aerosols and create better cloud models.

Every child, watching white fleecy shapes shift in the sky, naturally wonders where clouds come from. Most will probably learn that clouds form when water evaporates from lakes, rivers, and oceans and condenses into droplets that eventually fall as rain. Though mostly correct, that explanation is missing an important piece. Cloud droplets need airborne particles, called "aerosols," around which to condense. These particles are hurled into the sky naturally from dust storms, volcanoes, sea spray, and fires, but many come from human industrial and agricultural activities. That's right; those graceful, pristine puffs of cotton candy are also made of pollution and dirt that humans and the planet spew into the atmosphere.

Aerosols are tiny, ranging from 1/1000th of a micron to 100s of microns—anywhere between the size of a virus and a grain of sand. The period at the end of this sentence, at about 500 microns, would be an extremely large aerosol. Though miniscule, aerosols have an enormous effect on global climate by directly reflecting or absorbing solar radiation. Aerosols also indirectly affect the climate because they influence the size and abundance of cloud droplets, which in turn determine how much sunlight a cloud reflects and how much rain it produces.

Human-generated aerosols alter global climate, as do their better-known stepsisters, the greenhouse gases, but aerosols' effects are more complicated—so complicated that the Intergovernmental Panel on Climate Change (IPCC) declared them the greatest source of uncertainty in predicting climate change. And that's a big problem.

Just as a Florida homeowner relies on the local weather forecast to know if it's time to board up the windows, policy makers around the world need long-term, global-climate predictions to

Left: A wing-mounted probe collects aerosol samples over Alaska. Above: Inside the aircraft, Manvendra Dubey analyzes the samples using a photoacoustic instrument.

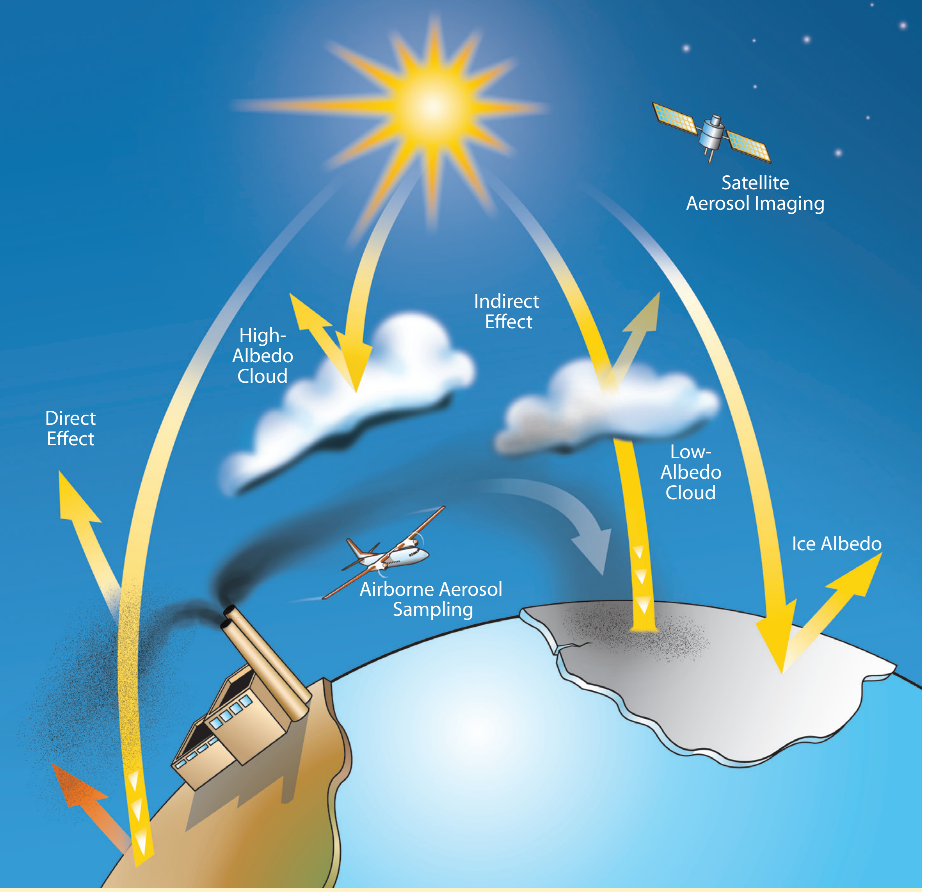

Climatic Effects of Aerosols

Aerosols, tiny particles suspended in the air, affect the Earth's energy balance by absorbing or reflecting solar radiation. This is called the aerosol direct effect. Aerosols also have an indirect effect by influencing cloud properties, such as lifetime, height, and the number and size of droplets. Clouds with fewer but larger droplets reflect less solar radiation (low albedo) while those with high aerosol concentrations reflect up to 90 percent of radiation back into space (high albedo). Dark aerosols from industry can settle on Arctic ice, which normally reflects sunlight, and accelerate melting by absorbing sunlight. Los Alamos researchers are analyzing aerosols using airborne and satellite techniques.

plan appropriate responses. Current climate models predict that by the end of this century there will be a temperature increase of 1.2°C to 4.4°C—a range far too wide to be a useful prediction. Aerosols are the unknown quantity.

To reduce the uncertainty, climate scientists must do two things: first, identify a precise numerical value for the effect aerosols have on the planet's energy budget. The Earth maintains a balance between incoming solar radiation and amounts reflected back into space or absorbed and re-radiated as infrared. Changes in the balance are known as "radiative forcing." Earth's preindustrial equilibrium is the baseline, and scientists need to accurately measure changes to it caused by human-generated aerosols. Second, scientists must gain a better understanding of the interactions between aerosols and clouds to accurately represent them in predictive models. To achieve these goals, researchers at Los Alamos National Laboratory are using the Laboratory's unparalleled computing power and advanced tools for gathering and interpreting atmospheric data.

Climate's Gray Area

We know that greenhouse gases, like CO_2, warm the planet, and we even know how much, but what about aerosols? "When you burn fossil fuels," explains Manvendra Dubey of Los Alamos' Earth and Environmental Sciences Division, "you emit not only CO_2, but also sulfur dioxide (SO_2), which becomes sulfate aerosols. Sulfates, along with some other aerosols, cool the planet by reflecting sunlight away." In the 1970s and 80s, before smokestacks were equipped to scrub sulfur, industry produced so much sulfate pollution that the resultant cooling counteracted the warming of greenhouse gases. But health concerns over particulate pollution and acid rain brought about the Clean Air Act, which forced industry to reduce sulfate emissions and ironically allowed the greenhouse effect to intensify. "Because we succeeded in reducing sulfate pollution," says Dubey, "we must now work twice as hard to control CO_2."

So if, as we learned with the sulfates, aerosols cool the planet, all we need to do is figure out how much and plug that number into climate models, right? Unfortunately, the problem is not so black and white—literally. Some aerosols, like sulfates and sea salt particles, may cool because they're white and reflect sunlight, but other aerosols, such as black carbon (soot), are dark and absorb sunlight like a black shirt on a hot summer day. When they all mix into one giant atmospheric pointillist painting, you get shades of gray whose effects are hard to quantify. But scientists are working to measure the amount of light reflected and absorbed by the aerosol mix—what they call "aerosol optical depth" (AOD). AOD is the key parameter in determining the elusive radiative-forcing number needed for precise climate predictions.

Listing to Dust

Conventional methods of measuring AOD can exaggerate the darkness by a factor of 2 or 3, but Dubey has corrected that error by deploying the world's first aircraft-mounted, three-laser, photoacoustic instrument. As a plane flies through clouds and haze, the instrument sucks in particles and exposes them to light from red, green, and blue lasers that together represent the solar spectrum. When particles absorb the laser energy and heat up, they expand the air around them and create a sound wave that is detected by a highly sensitive microphone. Sensors also detect light reflected by the particles. The two measurements together translate into an accurate measurement of AOD.

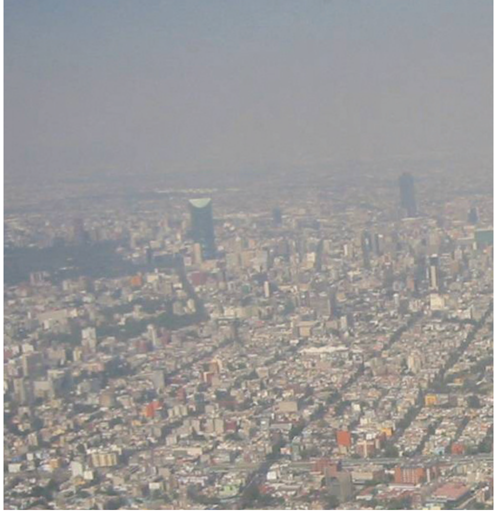

Soot and other aerosol pollutants rise into the atmosphere above Mexico City.
PHOTO BY NANCY MARLEY, ARGONNE NATIONAL LABORATORY

Dubey has flown over Mexico City, Korea, Houston, California, and even the Arctic to collect and analyze aerosol mixtures with the laser instrument. In Mexico City, he measured the effects of megacity pollution on global warming. In addition to finding sulfates and other kinds of aerosols he anticipated, he also detected a significant amount of aerosols produced by organic gases that vehicles emit. When the sun rises, these gas molecules undergo a photochemical reaction that turns them into particles dubbed "secondary organic aerosols." They were considered to be negligible, but climate modelers are now including them in chemical-transport simulations.

In Jeju, an island off the coast of South Korea and downwind from China, Dubey analyzed aerosols blowing in from Beijing to see if China had taken effective steps to clean the air for the 2008 Olympics. The Chinese government would not allow soot measurements within their country, but wind ignores borders, so Jeju was the next best thing.

But perhaps the most interesting of Dubey's observations were those obtained in the Arctic because they threw a new twist into the aerosol story. "You expect the Arctic to be pristine," says Dubey, "but it's pretty polluted." And the pollution isn't from local particles, as in Mexico City, but from an international mix of junk from all over the Northern Hemisphere. Dubey observed plumes similar to Los Angeles smog coming from Siberian fires, Gobi Desert dust storms, and industrial emissions. The new twist in the story is that the gray mix of aerosols, which contains a lot of black carbon, affects not just the atmosphere but also the ice itself. The aerosols settle on the Arctic ice sheet, causing it to absorb solar radiation and melt faster than computer models have predicted. Clean Arctic ice normally has a cooling, "albedo" effect, reflecting solar radiation.

A closer look at radiative forcing helps illustrate the significance of this phenomenon. Radiative forcing is expressed in watts per square meter (W/m^2). If a climatic influence warms the planet, as greenhouse gases do, it causes positive forcing. If it cools, as do sulfate aerosols, it causes negative forcing. Over the last century, human-caused greenhouse gases have produced a positive radiative forcing of 2.6 W/m^2 while aerosols are estimated to have had a negative radiative forcing of –1.2 W/m^2, though that figure is still highly uncertain. These measurements are global, long-term averages that don't take regional and seasonal effects into account. Those are the very effects that current ice-melt models lack. During the period Dubey studied the Arctic, the radiative forcing of black carbon for that region was a whopping 30 W/m^2.

"It's a double whammy," says Dubey. "Black carbon takes away the negative forcing of ice albedo and adds positive forcing directly to the ice surface." And it couldn't happen in a worse place.

In some places, regional and temporary disturbances might not impact global climate, but the Arctic isn't one of those places. Petr Chylek, a pioneer in aerosol science and a frequent collaborator with Dubey, explains why: "If global warming occurs, what disaster awaits humankind? Temperatures rose 0.7°C over the last 125 years, but you can't feel it. If it goes up another degree here, nothing happens. The danger is in the Arctic because melting Arctic ice results in rising sea levels. If the Greenland ice sheet melts, we have a global disaster."

Watching Aerosols from Space

Chylek was thinking about how aerosols affected climate long before it was a hot topic. Thirty-four years ago, while a researcher at the National Center for Atmospheric Research, Chylek published a paper entitled "Aerosols and Climate" in the prestigious journal *Science*. In that 1974 paper, he pointed out the need for accurate aerosol measurements and ended on a note of hope, speculating that "someday their effect may be measured directly when changes in the albedo of the earth-atmosphere system are remotely monitored by satellites." That day has come and Chylek, now remote-sensing team leader in Los Alamos' International, Space, and Response Division, is now in the satellite business.

Petr Chylek

Chylek's work builds on Los Alamos' history of using satellites to detect the illicit production of weapons of mass destruction. One of those satellites, the Multispectral Thermal Imager (MTI), has circled the planet since March 2000, collecting images of the Earth with instruments that see changes in light and heat that the human eye cannot. While looking for the telltale gases, dusts, and heat produced by chemical or nuclear activity, the MTI has produced mountains of environmental data. Chylek took on the task of figuring out what nondefense questions might also be answered with MTI's data, and there was the answer to his hopes from 1974—global pictures of aerosols.

But the MTI didn't just hand over ready-made answers. It provided measurements of radiance (light intensity) that Chylek and his team had to translate into accurate information about aerosol optical depth. After extensive calibrations with ground observations, Chylek hit upon a method that turned out to be highly accurate. For more global atmospheric data, Chylek turned to NASA's satellite-mounted MODIS instrument, which images the entire Earth every 1 to 2 days. While NASA had its own way to calculate aerosol optical depth, Chylek's method reduced the error by a factor of 2 to 3. This led to surprising observations of aerosol behavior over the Indian Ocean, a long way from the Arctic but involving the same culprit: black carbon.

During the winter, pollution increases dramatically over the Indian subcontinent, affecting cloud formation. As expected, Chylek found that more aerosols produced more and smaller cloud droplets, reflecting more sunlight back into space. But at altitudes where icy cirrus clouds form, the result was, surprisingly, the opposite—fewer but larger ice crystals letting more solar radiation through to the surface. This puzzling reaction was caused by the water-attracting properties of black carbon, and between the recent growth of industry and the longstanding practice of burning wood and coal as household fuels, India produces a lot of soot. In the dance between aerosols and water vapor, water snubs the other aerosols, and clings to the soot, which has the additional power to initiate freezing at just a few degrees below zero. In contrast, water condensed around sulfate particles can remain liquid up to –40°C. So water freezes quickly onto soot, depletes atmospheric moisture, and leaves many aerosols without a condensation partner. Climate models were not including this reaction.

Chylek's and Dubey's work complement each other well. Dubey measures aerosols over a city here and a region there while flying at different altitudes, which leads to highly detailed results, but it would require years to combine those details into a global picture. Chylek's satellite observations lack the detail, seeing

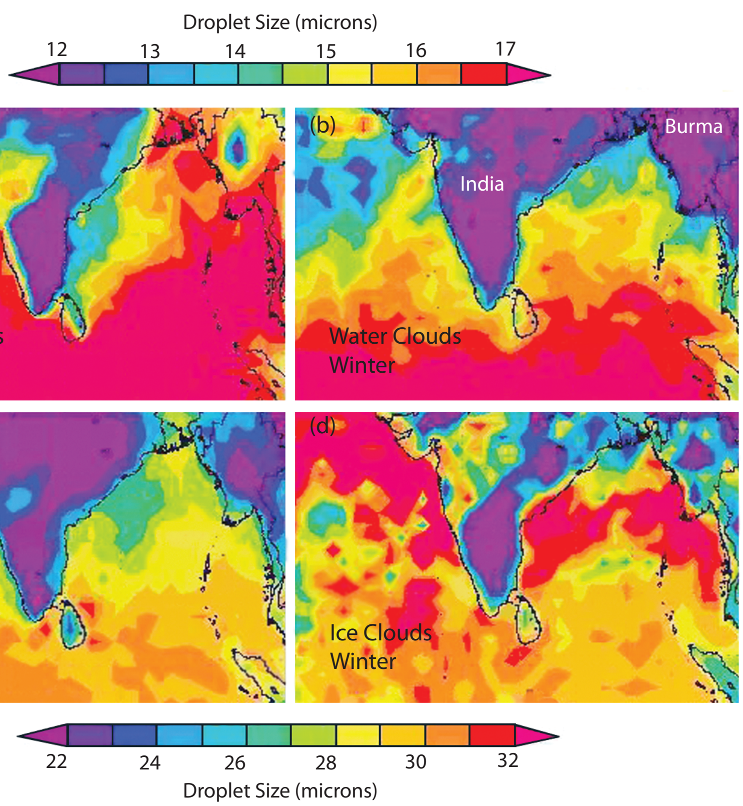

Satellite observations near India were as expected for lower-altitude water clouds but surprising for higher-altitude ice clouds. Increased pollution during the winter created many smaller cloud droplets (purple) in water clouds. But human-generated soot led to fewer but larger droplets in ice clouds, allowing more solar radiation through.

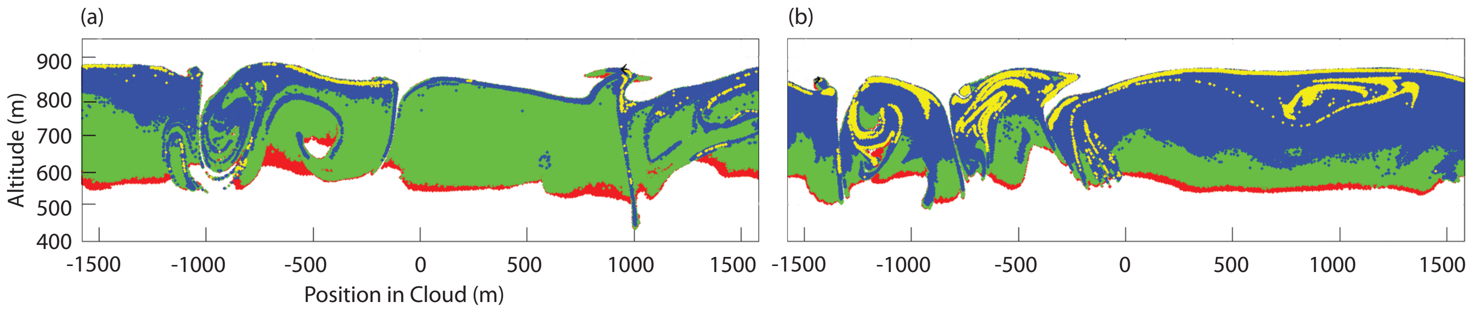

only optical averages of a whole column of aerosols from surface to satellite, but allow analysis of the entire planet in just days. Together, their work produces the stuff that better climate models are made of.

Cloud Modeling

Jon Reisner sits in a bare, modest office inside a small, weather-beaten, prefab building, but he's connected to one of the most-sophisticated supercomputing architectures on the planet. As he explains the complexities of cloud modeling, he keeps an eye on line after line of code from a lightning simulation scrolling down one of his oversized computer screens.

Reisner, of Los Alamos' Earth and Environmental Sciences Division, has made recent breakthroughs in accurately modeling aerosol-cloud processes. His success lies in his willingness to use an approach assumed unsuitable by most of the atmospheric modeling community. It's called the Lagrangian method. As Reisner worked through the physics and math, he found that the common objections to the method were unfounded.

Modeling aerosols and clouds is chiefly about simulating particle behavior, something Los Alamos has been perfecting since its founding. From the protons and neutrons of nuclear reactions to the toxic chemical and biological plumes of potential terrorist attacks, if it involves particles, Los Alamos has probably modeled it. Of the two common computational approaches to particle modeling, the Lagrangian and the Eulerian, the latter has been favored by atmospheric modelers, and that fact has led to some problems in simulating aerosol-cloud interactions.

The mathematical differences between the Lagrangian and Eulerian methods are not easily translated into words, but to get a simplified idea, think of a two-dimensional grid, like a checkerboard, with checkers representing aerosol particles. Modelers break large problems into grids to make complex calculations of the whole problem easier. In a Eulerian simulation, if a checker moves out of its grid square to cross the corner where four squares meet, it is no longer considered a single checker but a fraction of a checker in each of the four grid squares it partially covers. The checker—aerosol particle—gets diffused, which results in problems within computer simulations, especially at a cloud's edge, which is a border between droplets and no droplets. This Eulerian edge problem can cause clouds to instantly vanish from a simulation when, in reality, they would have survived a day or two.

In contrast, the Lagrangian approach solves the edge problem by tracking the checker/particle across grid lines, always representing it as a single undiffused checker. This results in more-accurate tracking and location of particles and also enables better calculations of particle collisions, which are important because colliding cloud droplets merge to form large, falling raindrops—a process that is not yet well understood.

Reisner's method was so accurate that Dubey turned to him to model aircraft observations of the effects of soot pollution on clouds. After 20 other models had failed to simulate the data, Reisner's method successfully re-created the observed response of clouds to soot.

An Emerging National Security Mission

A precise understanding of aerosols will take many more years of scientific work, but the good news is that aerosols are very short-lived compared with greenhouse gases. If we stopped pumping CO_2 into the atmosphere right now, most of what's already there would linger for another 50 to 200 years and some for thousands. But aerosols stay aloft from a few minutes to 10 days. Once we better understand their climatic effects and decide on controls, we can deal with aerosols quickly—if everyone cooperates.

Earth's climate is an international problem, and international solutions require treaties. Treaties work only if we can verify compliance. Detecting and tracking man-made aerosols and greenhouse gases against the background noise of natural emissions is no easy task, but it's the same kind of task that Dubey, Chylek, Reisner, and others are undertaking to measure and model aerosols. The knowledge honed through their research will produce the expertise and tools needed to verify compliance with environmental treaties. Such treaties, like the Nuclear Test Ban Treaty of the past, may be integral parts of the future national security landscape.❖

—Anthony Mancino

Above: Cloud simulations made using a Lagrangian modeling approach. Red indicates the smallest droplets, followed by green, blue, and the largest, yellow. High aerosol concentrations (left) lead to many small droplets while lower concentrations (right) produce fewer larger droplets.

DIALOGUE

Taking Charge

A Discussion about Electrical Energy Storage

The Office of Science at the Department of Energy (DOE) is looking for revolutionary breakthroughs in electrical energy storage, describing them as "perhaps the most crucial need for this nation's secure energy future." Laboratory scientists Albert Migliori and David Thorn discuss the need for energy storage and how Los Alamos can help meet it.

1663: We often hear about the need to conserve energy or to find alternative energy sources but rarely about the need to store it. Why does the DOE consider electrical energy storage "crucial" to our secure energy future?

Thorn: The answer is really that the world uses about 17 *trillion* joules of energy per second, and that rate is likely to double by 2050. We're going to need energy and lots of it. Simply burning more coal and oil isn't a solution, partly because oil supplies are dwindling, but mostly because the carbon dioxide that's emitted when you use those sources drives global climate change. To meet the impending energy need, we'll have to ramp up the use of carbon-free energy sources such as nuclear, wind, and solar.

Electrical energy storage comes into play because wind and solar generate electricity intermittently—we can't control when or how much they supply. Now Albert will tell you that because of the conservation of energy, the electrical supply must always meet the instantaneous demand. If it doesn't, even for just a few seconds, large portions of the power grid can go down. So by increasing the use of wind and solar power, we introduce a degree of uncertainty into our electric supply and actually make the grid less stable. Electrical energy storage lets you save electricity for later use, so it mitigates wind and solar's inherent variability. It enables the expansion of those two energy sources, since now they can be added to the grid without de-stabilizing it.

Migliori: Consider how New Mexico's public utility, PNM, sees it. Suppose the wind really kicks up in the dead of night, when the demand for electricity is low and is already being met. They can't put that electricity on the grid, so if there's no place to store it, they have to shut the wind turbines down. Not only is that a waste of a capital investment, but the wind farm is idle when it could be making more power than a coal-fired power plant.

Conversely, in Albuquerque at about 4 p.m. during the summer, many air conditioners come on as people get home from work and PNM has to deal with a huge surge in electric usage. Solar power would help, but the sun is waning, so you're well off the peak for generating solar electricity. Unfortunately, in New Mexico, the strong winds blow in the spring at night, not in the summer in the afternoon, so the wind farm doesn't help meet that demand either. With storage, you can collect solar energy during the day and then use it at night or store the spring winds and offset the summer's air conditioning energy needs.

1663: Store the spring winds? Isn't that a lot of energy?

Migliori: Yes. There are minutes in New Mexico when 25 percent of our power comes from wind farms—over 200 million watts from wind power! It's the highest fractional utilization of any state. There's enough energy generated by New Mexico wind farms in March and April to run a utility grid for weeks without turning on a coal or nuclear plant.

Thorn: The problem is we don't have the means to store that much electrical energy, and conservative, incremental improvements in storage technology won't get us there. We need revolutionary changes.

Migliori: But even if you develop a high-tech battery or something that could store two-months worth of wind energy,

To download the 2007 DOE Office of Science report on electrical energy storage, go online to http://www.sc.doe.gov/bes/reports/files/EES_rpt.pdf

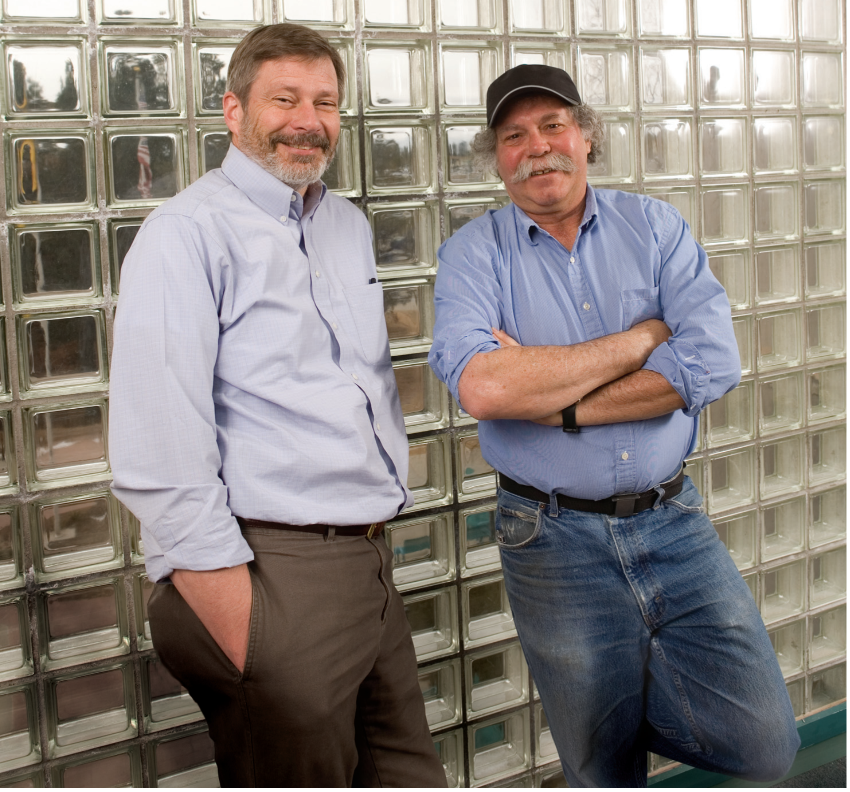

David Thorn (left) and Albert Migliori

all-that energy would be parked on a windswept ridge miles from anything. To get it to Albuquerque, you'd have to build a costly, high-power transmission line.

The idea is to go to a distributed storage system. This is something the Laboratory has looked into and is ready to contribute to in a major way.

There's already a utility grid that carries power to every New Mexico home. What if every house had a little bit of electrical-storage capacity, say a bank of lithium-ion batteries that sat in a box beneath the electric meter? You could then use the existing grid to distribute the output of your wind farm to thousands of these storage boxes. There would be no need to build new transmission lines.

Plus, with distributed storage, a state or regional utility could run its coal-fired or nuclear power plants all the time under optimum conditions so that the plants generated the minimum amount of atmospheric carbon or nuclear waste per unit of electricity. The utility would charge up these storage boxes when there was power to burn, so to speak, then recover the electricity when needed.

Ultimately, the utility would have some control over when electricity went to, say, the boiler, air conditioner, and dishwasher in a house. Then during the summer, they could stagger air conditioner "on" times, even if it was just for 10 minutes or so, to allow time for more power sources to come online and supply electricity.

Thorn: A similar idea is being floated right now because once you put that kind of communication and control in place, a utility could really optimize the use of its resources.

But the grid is a massive network, and distributed storage entails integrating thousands of storage units into it. Making it work efficiently would involve a significant modeling and computation effort that would be just as important as developing the electrical-storage units themselves.

Migliori: With our expertise in both modeling and computation, that problem is tailor-made for Los Alamos. It's a wonderful opportunity.

Thorn: We also have to consider that in the United States, about one-third of our energy goes towards heating homes and buildings, one-third goes towards producing electricity, and one-third goes towards transportation. Because transportation is strongly dependent on foreign oil, there's a great motivation, maybe even an imperative, to develop alternative fuels and/or electrical storage systems that could be used to power an electric car.

1663: Transportation is a different storage need?

Thorn: Yes. The main obstacle there is the energy density of the storage medium. We're used to gasoline, which has a remarkably high energy density of about 11 kilocalories per

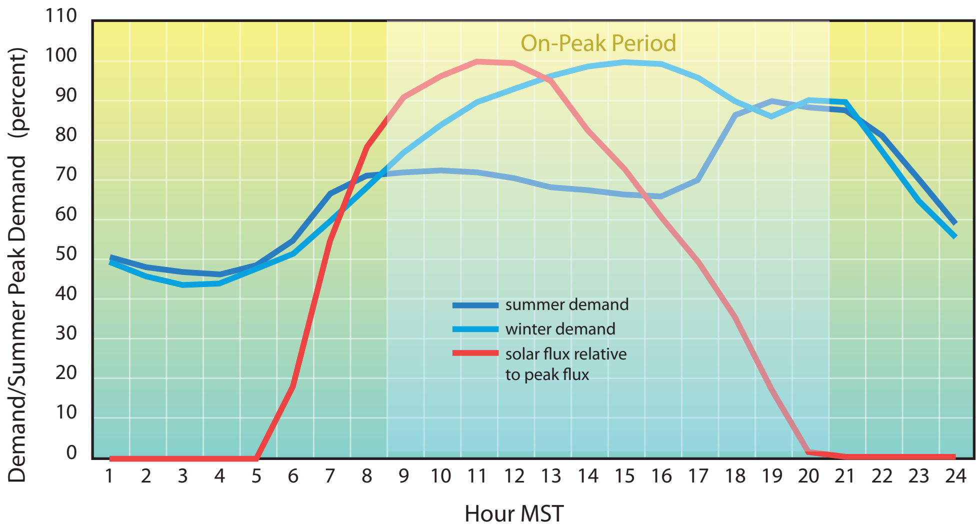

In New Mexico, the amount of electricity generated by a solar panel peaks around noon. (red curve), whereas the demand for electricity peaks closer to 4 p.m. during the summer, 7 p.m. during the winter. Storing the solar-generated electricity would allow the utility to meet demand.

Wind power accounts for about 8 percent of New Mexico's electricity, but only when the wind blows. Some type of electrical energy storage will be needed to allow the large-scale use of wind power.

gram (kcal/gm). When you burn an entire tank of gasoline, you liberate about 250,000 kcal of energy—a billion joules!

Migliori: Remember that only about 20 percent of that energy is used to move the car. The rest is lost in the form of heat and through inefficiencies in the car's power train. An electric car would have a much-higher efficiently.

Thorn: So given what Al said, and considering that the best lithium-ion battery has an energy density of at most 0.4 kcal/gm, we've got to develop an electrical storage device that has roughly five times the energy density of our best battery if we want an electric car that performs as well as a gasoline-powered vehicle. We don't know how to do that yet, but I'll wager it will happen not through engineering but through an understanding of fundamental science.

1663: What about using hydrogen?

Thorn: Hydrogen is an energy storage medium. A common scenario is to use electricity from a power plant to produce hydrogen through water electrolysis. Two electrodes are placed in water. When a current is run between them, the electrical energy overwhelms the chemical bonds, and the water separates into hydrogen and oxygen. You store the gases separately, then when needed, feed the hydrogen into a fuel cell to generate electricity.

Migliori: Hydrogen gas has such a low density that it's hard to imagine its becoming a commodity the way oil or gasoline is. For example, to carry enough hydrogen to power a car for 200 miles, you'd have to load it into a "bottle" at extremely high pressure, which raises all kinds of flags about safety, refueling, delivery, etc. So people have been looking at ways to store hydrogen not as a gas but as part of a molecular compound.

Thorn: Currently, one of our best hopes is ammonia borane, a relatively dense solid at room temperature that's about 18 percent hydrogen by weight. You chemically remove the hydrogen to run your fuel cell. Los Alamos established collaborations for working on ways to improve the efficiency and lower the cost of regenerating the spent material. We're also heavily engaged in research to make fuel cells less expensive, more robust, and more accommodating of fuels other than hydrogen. And across the Laboratory, scientists and engineers are working to improve power plant designs and researching grid stabilization, nuclear power, etc.

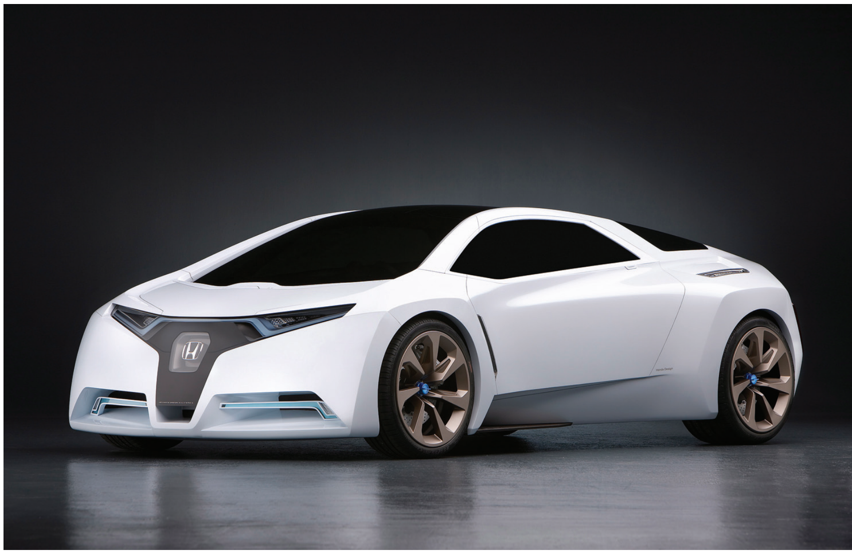

COURTESY OF HONDA

1663: Suppose a hydrogen-powered car becomes reality. How will that affect the oil industry and the economy?

Thorn: Any changes to our current energy paradigm will have far-reaching consequences. There's no simple, single answer. We have an Energy Security Center here at the Laboratory that focuses on those questions.

Migliori: Getting back to what Dave said earlier about revolutionary change coming from basic science, I believe we must understand the fundamental electrochemical processes that occur at the electrode surfaces in batteries and other storage devices.

Above: This concept car is powered by an electric engine and "fueled" by hydrogen gas. The hydrogen is a form of electrical energy storage. Electricity from a power plant generates the hydrogen by water electrolysis. The hydrogen, stored in two high-pressure tanks in the car, runs through an on-board fuel cell to regenerate the electricity.

Consider this. During electrolysis, electric fields develop in nanometer-size regions of the electrode surface—regions that contain less than a few hundred atoms—that have strengths on the order of 10 billion volts per meter. That's about 10,000 times the electric field in a lighting bolt. We're talking enormous energies, enormous electric fields. But after 250 years of electrochemisty, we still don't know exactly what these nanoscale fields look like or how they behave. All of our equations fail miserably at the nanoscale.

Thorn: That's partly because we haven't had the tools to measure the electric field at that scale. You can't just stick a voltmeter in there. Most techniques probe much-larger areas, so you end up measuring the properties of the surroundings and not of the tiny nanoscale interface.

Migliori: But the wonderful thing about nanotechnology is that we can carve up a silicon wafer, coat it however we want with one or several layers of metal atoms, and make thousands of identical nanoscale test devices. So we can get a greatly enhanced signal and know that it's coming from the structures we're interested in.

With the Center for Integrated Nanotechnology right here at Los Alamos, plus a strong collaboration with the College of Nanoscience and Engineering in Albany, New York, we can make the nanostructures needed to do this research. Couple that with advanced measurement techniques and the most-sophisticated computational tools in the world, and we can begin to acquire knowledge about the most-fundamental electrochemical processes and use that knowledge to build a better theoretical model.

Thorn: From a device scenario, if you had a lithium battery based on huge numbers of nano entities fabricated the right way, they would probably fail one at a time in a way that you could probably tolerate. So the actual construct will offer advantages to doing things in bulk.

Migliori: And from a research perspective—you know the old saw that you can't have an electric field inside a metal? Not true. But to eliminate that electric field, the free electrons in the metal have to migrate and form a surface layer of charge, which creates an electric field that effectively cancels any field in the metal.

At a nanoscale interface—maybe a hundred metal atoms sitting on a silicon substrate—there aren't enough electrons to cancel the field. Instead, the electrons will be pushed around by any external electric field and pressed into a smaller volume. When that happens, the Heisenberg uncertainty principle says that the kinetic energy of the electrons goes up.

In fact, that additional kinetic energy becomes a significant fraction of the energy stored. This is called "quantum capacitance." All of a sudden you've taken a potential-energy storage problem and introduced a component that's a kinetic-energy storage problem. What a fantastic thing for the Lab to be working on, understanding the details of that.

Thorn: Still, the challenge as I see it is how to devise nanostructures that are different enough from the bulk materials to give us large gains in energy densities. The possibility is there, but we don't understand how to do it.

Migliori: I believe it will happen. Our job is to work on gaining a fundamental understanding of the juncture between physics, chemistry, and nanotechnology. If we do our job, others will have a chance to develop really good revolutionary storage devices, and the game will be over. ❖

—Jay Schecker

With the appropriate type of storage, solar energy could potentially supply all of the nation's electricity.

SPOTLIGHT

A Sound Overhaul for Flow Cytometry

Have you heard? A Laboratory-developed innovation is using sound waves to make cells or particles line up for accurate, fast analysis. The new technology—acoustic flow cytometry—may bring much-needed medical diagnostics to many parts of the Third World.

Flow cytometry is particularly useful to the medical profession, which uses flow cytometers for producing blood counts and monitoring the progress of HIV/AIDS patients. The standard instrument uses a hydrodynamics (fluidics) system to transport cells, one by one, through a horizontally focused laser beam.The cells are suspended in saline solution that flows at the center of, and in the same direction as, a larger and faster-moving stream (a "fluid sheath"). This configuration forces the cells into single file for their trip through the laser beam, so the light they scatter or emit can be readily detected and analyzed by a data system.

But the bulk, expense, and fragility of the hydrodynamics system have long confined cytometers to the most-sophisticated laboratories and clinics.

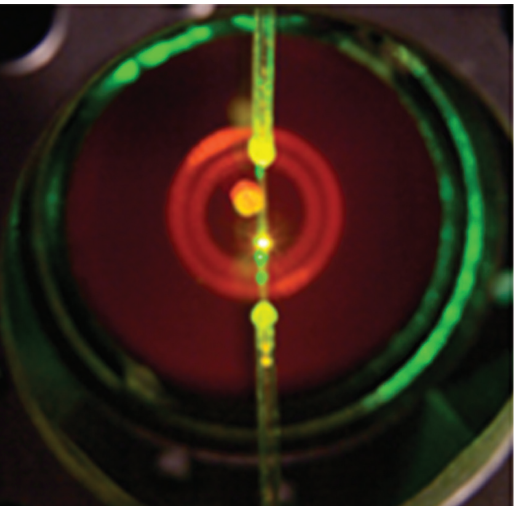

The view down the new cytometer's glass capillary (perimeter in red), reveals a particle (yellow) passing through the laser beam.

The new technology, licensed by the Laboratory to Acoustic Cytometry Systems (ACS), a Los Alamos company, eliminates the fluid sheath. Instead, it sends the suspended cells through a glass capillary that has a "piezoceramic" acoustic source bonded to its outer wall. The source converts an electric charge into an ultrasonic field that drives the cells to the center of the capillary.

The new design makes cytometers smaller, lighter, less complicated, and also less expensive. Compared with the fluid sheath, the ultrasonic field produces a slower transit of cells through the laser. Because each cell gets interrogated longer, scattering and emitting more light, a smaller, less-expensive laser and less-demanding optics can be used.

All these advantages—including eliminating the sheath's purified water, which may be rare in Third-World countries— may free flow cytometry for wider use. Indeed, acoustic cytometry is about to go global. ACS has been by purchased by California's Invitrogen Corporation, which provides life-science products and services around the world. So a Laboratory-developed technology will soon be helping many more people.

Quantum Slip

Air-bag actuators, handhelds such as iPhones, and video-game controllers all get motion or orientation information from tiny motion sensors whose even-tinier parts rotate or bend freely when the device is spun or accelerated.

As Diego Dalvit of T-Division knows, an arcane quantum-mechanical force—the "Casimir" force—can stop those parts from rotating or bending at all. But he and his Los Alamos team appear to have found a way to keep things moving.

Each motion sensor is usually fabricated on a chip called a microelectromechanical system, or MEMS. As MEMS are made smaller, several forces arise that can cause the sensor's moveable parts to stick to nearby surfaces. For nanoscale MEMS—devices a thousand times smaller than the width of a human hair—the Casimir force dominates. In fact, unless it can be neutralized, the Casimir force threatens to halt the progress of the incredible shrinking MEMS.

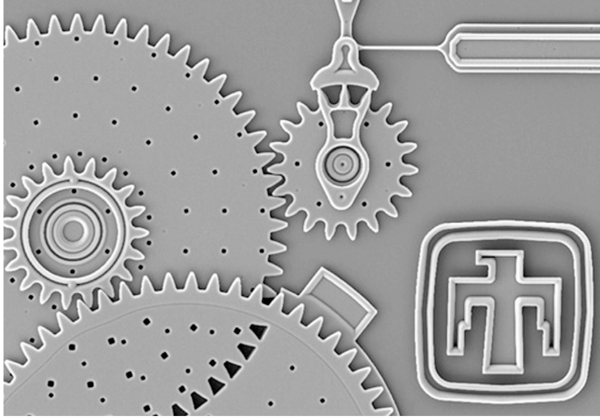

Microgears fabricated at Sandia National Laboratories. As movable parts get smaller, the sticky Casimir force can make them grind to a halt.
COURTESY OF SANDIA NATIONAL LABORATORIES

The Casimir force is subtle. Quantum physics predicts that photons can suddenly appear and disappear from the vacuum in a very short time. During their fleeting existence, these "virtual" photons exert a "radiation" pressure on surfaces, in the same way that sunlight pushes on comet tails. For example, between two thin parallel conducting plates, the only wavelengths of light that can exist are those that exactly match the distance between the opposing surfaces of the plates. Outside the plates, the light has no such constraints. As a result, the radiation pressure of the virtual photons outside the plates is greater than it is between them, so the plates are pushed together.

Previous theoretical studies of the Casimir force considered only conducting or semiconducting surfaces, for which the force is always attractive. However, Dalvit's team found that special magnetic metamaterials—materials whose properties derive from tiny structures patterned onto their surface—can neutralize the attraction and even make the Casimir force repulsive! Team members are planning an experiment to test the theory. If it's right, then nanoscale MEMS with nonstick, metamaterial parts could make for some very freewheeling devices.

From Homeland Security to High-Performance Fabric

A microelectonics revolution arose from the use of plasma for manufacturing computer chips. Now plasma, an ionized gas consisting of electrons, ions, and chemically active neutral fragments, is revolutionizing the clothes we wear.

Former Los Alamos staff member Gary Selwyn recognized that plasma used in a vacuum (as it is for computer chips) might also be used at atmospheric pressure, under the right conditions. Pursuing that idea, he developed a technique in which a jet of plasma was used to decontaminate vehicles and equipment. It was, in effect, one of the Laboratory's first homeland security projects. To commercialize the plasma jet, Selwyn founded APJeT, a Santa Fe company, and licensed the technology from the Laboratory.

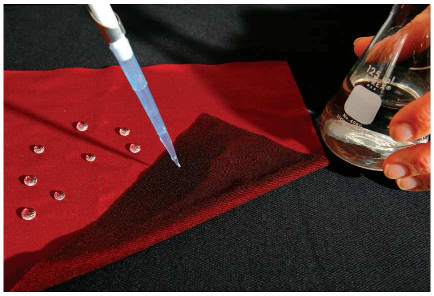

The new fabric treatment can produce water resistance on one side (see water droplets) and wicking on the flip side.

The high cost of vacuum plasma had kept plasma treatment from being used for commodity items, but the new atmospheric plasma-jet opened the door. APJeT targeted fabric treatment for the high-performance outdoor clothing market.

Using helium for the plasma gas was one key to success, allowing APJeT to develop a high-density (for treatment speed) atmospheric plasma that was also "nonthermal"—cool enough for use with fabric. Helium has a high thermal conductivity that allows any heat generated to be easily removed. In addition, helium is unique in that it prevents arcing (sparking), a common problem with atmospheric plasmas. Suddenly, the complicated, vacuum-based plasma that was used before could be created at atmospheric pressure and room temperature.

As a result, APJeT's methods can take a commodity textile product—woven or knitted polyester—and turn it *at low cost* into a product rivaling DuPont and Gore-Tex products. In addition, because APJeT's fabric finishing is plasma-based and so done in the gas phase (unlike traditional "wet" methods), fabrics can be given sequential single-side treatments to produce a product that repels water and stains on the outside and absorbs moisture on the inside.

APJeT's plasma machine is manufactured under license by Morrison Plasma Systems. Installed at the College of Textiles at North Carolina State University, it is currently being demonstrated for customers.

HIV's Evolving Evolution

The human immunodeficiency virus (HIV)—the virus that causes AIDS—is such a tenacious pathogen because its genetic material readily mutates when the virus reproduces. Once inside the body, the HIV population quickly evolves and diversifies into thousands of different viruses, so-called quasi-species, making it nearly certain that some viruses will evade the body's immune system and resist antiviral drugs.

Curiously, during the later stages of HIV infection, the quasi-species begin to accumulate fewer new genetic differences from the original viral strain (the divergence of genetic sequences begins to saturate), and the number of different viruses (the viral diversity) begins to decline. These changes always precede the progression of the disease to AIDS, although the time for AIDS to develop varies greatly from patient to patient.

Understanding the dynamics of these changes could help scientists develop new ways to stop the virus. That was why Los Alamos researchers Ha Youn Lee, Alan S. Perelson, and Thomas Leitner, all of the Theoretical Biology and Biophysics group, and collaborator Su-Chan Park from Cologne University, Germany, developed a simple model of HIV sequence evolution. The model had two main components: (1) fitness, the number of offspring produced, and (2) the proportion of offspring that are mutants. They tested the model using data from the Los Alamos HIV Sequence Database, which holds more than 250,000 genetic sequences of HIV from patients around the globe.

The results were surprising. In short, fitting the model to data showed that after evolving at a constant rate, the quasi-species' rate of evolution slowed down.

HIV infects cells of the immune system, reproduces within them, and then kills the cells, a course of action that suggests several possible reasons for the slowdown. One possibility is that because there are fewer immune cells in the later stages of the infection, the reproduction rate decreases. Another is that the weakened immune system is no longer able to apply "selective pressure" to drive the evolution. In any case, the work has already reconciled previously conflicting observations of the relationships between the rate of HIV evolution and disease progression.

The work is reported in a recent paper, "Dynamic Correlation between Intrahost HIV-1 Quasispecies Evolution and Disease Progression," *PLoS Computational Biology* 4 (12), e1000240 (2008).

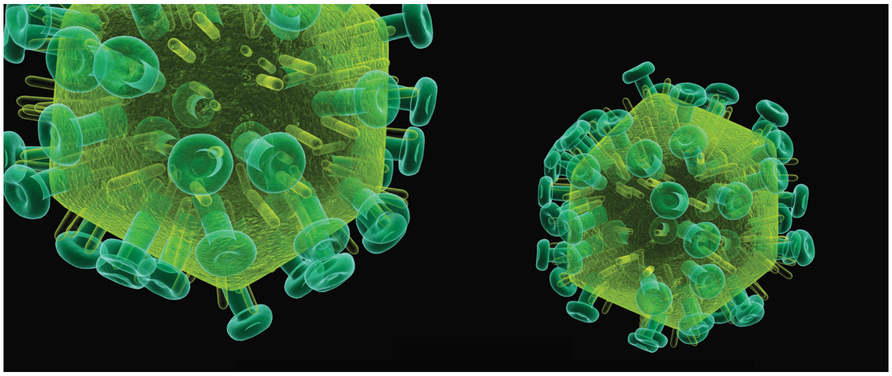

Artist's conception of HIV. Apparently, the virus's rate of evolution gradually slows.

Presorted Standard
U.S. Postage Paid
Albuquerque, NM
Permit No. 532

Or Current Resident

ISSN: 1942-6631

Address mail to
1663
Mail Stop M711
Los Alamos National Laboratory
Los Alamos, NM 87545
email: 1663magazine@lanl.gov
Tel: 505-667-1447
Fax: 505-665-4408
www.lanl.gov/science/1663/

Springtime in northern New Mexico

Los Alamos NATIONAL LABORATORY EST. 1943

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396.
A U.S. Department of Energy Laboratory LALP-09-028

Printed on recycled paper